arXiv:1403.2310v2 [stat.ME] 19 Aug 2015

Adaptive Penalized Estimation of Directed Acyclic Graphs From Categorical Data

Fei Fu*, Jiaying Gu*, and Qing Zhou†

Department of Statistics, University of California, Los Angeles, CA 90095, USA

Abstract

We develop in this article a penalized likelihood method to estimate sparse Bayesian networks from categorical data. The structure of a Bayesian network is represented by a directed acyclic graph (DAG). We model the conditional distribution of a node given its parents by the multi-logit regression and estimate the structure of a DAG via maximizing a regularized likelihood. The adaptive group Lasso penalty is employed to encourage sparsity by selecting grouped dummy variables encoding the level of a factor. We develop a blockwise coordinate descent algorithm to solve the penalized likelihood problem subject to the acyclicity constraint of a DAG. When intervention data are available, our method may construct a causal network, in which a directed edge represents a causal relation. We apply our method to various simulated networks and a real biological network. The results show that our method is very competitive, compared to other existing methods, in DAG estimation from both interventional and high-dimensional observational data. We also establish consistency in parameter and structure estimation for our method when the number of nodes is fixed.

KEY WORDS: Adaptive group Lasso; Coordinate descent; Discrete Bayesian network; Multi-logit regression; Structure learning.

*These two authors contributed equally to this work.

†To whom correspondence should be addressed (email: zhou@stat.ucla.edu). This work was supported by NSF grants DMS-1055286 and DMS-1308376 (to Q.Z.).

1 Introduction

A Bayesian network is a probabilistic graphical model whose structure encodes the conditional independence relationships among a set of random variables. It is graphically represented by a directed acyclic graph (DAG). Recent years have seen its popularity in the biological and medical sciences for inferring gene regulatory networks and cellular networks, partially attributed to the fact that it can be used for causal inference. Learning the structure of these biological networks from data is a key to understanding their functions. Most methods that have been proposed for structure learning of DAGs fall into two categories.

The first category encompasses the so-called constraint-based methods that rely on a set of conditional independence tests. The tests are used to examine the existence of edges between nodes. The idea of these methods is to find edges that cannot exist under certain assumptions. In practice the assumption can be so strong that it constitutes the main drawback of these methods. The PC algorithm proposed by Spirtes et al. (1993) and the MMPC algorithm proposed by Tsamardinos et al. (2006) are two well-known examples. The second category includes score-based methods whose goal is to search for a DAG that maximizes certain scoring function. The scoring functions that have been employed include several Bayesian Dirichlet metrics (Buntine 1991; Cooper and Herskovits 1992; Heckerman et al. 1995), Bayesian information criterion (Chickering and Heckerman 1997), minimum description length (Bouckaert 1993; Suzuki 1993; Bouckaert 1994; Lam and Bacchus 1994), entropy (Herskovits and Cooper 1990), et cetera. There are also Monte Carlo methods (Ellis and Wong 2008; Zhou 2011) which draw a sample of DAGs from a posterior distribution. Additionally, there are hybrid methods which combine the two approaches. The idea is to use a constraint-based method first to narrow the search space, and then use a score-based method to learn the DAG structure (Tsamardinos et al. 2006; Gámez et al. 2011).

With the rising interest in sparse statistical modeling, score-based methods seem particularly attractive since various sparse regularization techniques are potentially applicable to them. Assuming a given natural ordering among the nodes, Shojaie and Michailidis (2010) decomposed DAG estimation into a sequence of ℓ_1-penalized linear regression problems. Fu and Zhou (2013)

recently developed an ℓ_1-penalized likelihood approach to structure estimation of sparse DAGs from Gaussian data without assuming a given ordering. This method has been further generalized to the use of concave penalties by Aragam and Zhou (2015). Other recent developments on penalized DAG estimation include the work of Schmidt et al. (2007) and Xiang and Kim (2013). There is also theoretical development on ℓ_0-penalized estimation of sparse high-dimensional DAGs under a multivariate Gaussian model (van de Geer and Bühlmann 2013).

Despite the recent fast developments on sparse regularization methods for learning Gaussian DAGs, a generalization to categorical data is highly nontrivial. First, each node now represents a factor coded by a group of dummy variables. In order to select a group of dummy variables together, we need to use the group Lasso penalty instead of penalizing individual coefficients. Second, the log-likelihood for categorical data is more complicated than that for Gaussian data, and consequently, development of an algorithm to maximize the penalized log-likelihood becomes much more challenging. In this paper, we propose a principled generalization of the penalized likelihood methodology in our previous work (Fu and Zhou 2013) to estimate sparse DAGs from categorical data without knowing the ordering among variables. Our method may take both observational and interventional data. We demonstrate that our penalized estimator can recover a causal network from sufficient interventional data. Furthermore, we will present empirical evidence for the application of our method on high-dimensional observational data.

The remainder of this paper is organized as follows. Section 2 gives a brief introduction to discrete Bayesian networks and describes a multi-logit model that we propose to formulate the structure learning problem. Section 3 develops a blockwise coordinate descent algorithm used for parameter estimation, where each coordinate descent step is achieved by applying a quadratic approximation iteratively. Section 4 establishes asymptotic properties of our penalized estimator. Section 5 reports numerical results of our method on four types of simulated networks for both interventional and observational data, and Section 6 presents results on a biological data set. These two sections also include extensive comparisons with other competing methods. The paper is concluded with a discussion in Section 7. Technical proofs are relegated to the Appendix.

2 Problem Formulation

2.1 Discrete Bayesian networks

The structure of a Bayesian network with p random variables $X_1, \ldots, X_p$ is given by a DAG $\mathcal{G} = (V, E)$. The set of nodes $V = \{1, \ldots, p\}$ represents the set of random variables $\{X_1, \ldots, X_p\}$, and the set of edges $E = \{(j, i) \in V \times V : j \to i \text{ is a directed edge in } \mathcal{G}\}$. Given the structure of $\mathcal{G}$, the joint probability density (mass) function of $\{X_1, \ldots, X_p\}$ can be factorized as

$$p(x_1, \ldots, x_p) = \prod_{i=1}^{p} p(x_i | \Pi_i^{\mathcal{G}}), \tag{1}$$

where $\Pi_i^{\mathcal{G}} = \{j \in V : (j, i) \in E\}$ is called the set of parents of X_i and $p(x_i|\Pi_i^{\mathcal{G}})$ denotes the conditional density of X_i given $\Pi_i^{\mathcal{G}}$. Throughout the paper, we use j and X_j interchangeably.

Given a joint distribution, there may exist multiple factorizations of the form in (1), leading to different DAGs. The DAGs encoding the same set of joint distributions form an equivalence class of Bayesian networks. Therefore, it may not be possible to distinguish equivalent DAGs from observational data alone. See Aragam and Zhou (2015) for related discussion on Gaussian Bayesian networks. However, when used for causal inference, equivalent DAGs do not have the same causal interpretation and can be differentiated based on experimental data. There are methods for learning causal relations from a mix of observational and experimental data (Cooper and Yoo 1999; Meganck et al. 2006; Ellis and Wong 2008) and related work on inferring gene networks from perturbed expression data (Peér et al. 2001; Pournara and Wernisch 2004).

We describe below how the joint distribution of a Bayesian network can be modified to incorporate experimental data. For a detailed account of causal inference using Bayesian networks, please refer to Pearl (2000) and references therein. Briefly, assuming X_i, $i \in \mathcal{M} \subset \{1, \ldots, p\}$, is under experimental intervention, the joint density in (1) becomes

$$p(x_1, \ldots, x_p) = \prod_{i \notin \mathcal{M}} p(x_i | \Pi_i^{\mathcal{G}}) \prod_{i \in \mathcal{M}} p(x_i | \bullet), \tag{2}$$

where $p(x_i|\bullet)$ specifies the distribution of X_i under intervention. Experimental data generated

from $\mathcal{G}$ can therefore be considered as being generated from the DAG $\mathcal{G}'$ obtained by removing all directed edges in $\mathcal{G}$ pointing to the variables under intervention. It should be noted that (2) also applies to observational data for which $\mathcal{M}$ is simply empty. Hereafter, we develop our method under the assumption that part of the data are generated under experimental intervention, while regarding purely observational data as the special case of $\mathcal{M} = \varnothing$.

In a discrete Bayesian network, each variable X_i is considered a factor with r_i levels, indexed by $\{1, \ldots, r_i\}$. The set of its parents $\Pi_i^{\mathcal{G}}$ has a total of $q_i = \prod_{X_j \in \Pi_i^{\mathcal{G}}} r_j$ possible joint states $\{\boldsymbol{\pi}_k : k = 1, \ldots, q_i\}$. Let $\Theta_{ijk} = P\left(X_i = j \mid \Pi_i^{\mathcal{G}} = \boldsymbol{\pi}_k\right)$. A discrete Bayesian network $\mathcal{G}$ may be parameterized by $\boldsymbol{\Theta} = \{\Theta_{ijk} \geq 0 : \sum_j \Theta_{ijk} = 1\}$ via a product multinomial model given the graph structure. Denote by $\text{supp}(\mathbf{v})$ the set of nonzero components of a vector $\mathbf{v}$. The number of parameters in this product multinomial model

$$|\text{supp}(\boldsymbol{\Theta})| = \sum_{i=1}^{p} r_i q_i = \sum_{i=1}^{p} r_i \prod_{X_j \in \Pi_i^{\mathcal{G}}} r_j.$$

If we assume that every variable has the same number of levels, i.e., $r_i = r$ for all i, then

$$|\text{supp}(\boldsymbol{\Theta})| = \sum_{i=1}^{p} r^{1+|\Pi_i^{\mathcal{G}}|}, \tag{3}$$

which grows exponentially as the size of the parent set $|\Pi_i^{\mathcal{G}}|$ increases. To reduce the number of free parameters, we propose a multi-logit model for discrete Bayesian networks under which development of a penalized likelihood method is straightforward. For the same DAG structure, the number of parameters can be much smaller compared to the product multinomial model.

2.2 A multi-logit model

We encode the r_i levels of X_i, $i = 1, \ldots, p$, by a group of $d_i = r_i - 1$ dummy variables. Suppose that we have a data set $\mathcal{X} = (\mathcal{X}_{hi})_{n \times p}$ generated from a causal discrete Bayesian network $\mathcal{G}$, where $\mathcal{X}_{hi}$ is the level of X_i in the h^{th} data point, $h = 1, \ldots, n$. Let $\mathbf{x}_{h,i} \in \{0,1\}^{d_i}$ be the group of dummy variables for $\mathcal{X}_{hi}$ and $\mathbf{x}_h = (1, \mathbf{x}_{h,1}, \ldots, \mathbf{x}_{h,p}) \in \{0,1\}^d$, where $d = 1 + \sum_{i=1}^{p} d_i$. (Hereafter, all vectors are regarded as column vectors.) For a discrete Bayesian network $\mathcal{G}$, we

propose to model the conditional distribution $[X_j|\Pi_j^{\mathcal{G}}]$ using the following multi-logit model

$$\begin{aligned} p_{j\ell}(\mathbf{x}_h) &\triangleq P(X_j = \ell|\mathbf{x}_h) \\ &= \frac{\exp(\beta_{j\ell 0} + \sum_{i=1}^{p} \mathbf{x}_{h,i}^T \boldsymbol{\beta}_{j\ell i})}{\sum_{m=1}^{r_j} \exp(\beta_{jm0} + \sum_{i=1}^{p} \mathbf{x}_{h,i}^T \boldsymbol{\beta}_{jmi})} \\ &= \frac{\exp(\mathbf{x}_h^T \boldsymbol{\beta}_{j\ell\cdot})}{\sum_{m=1}^{r_j} \exp(\mathbf{x}_h^T \boldsymbol{\beta}_{jm\cdot})}, \ \ell = 1, \ldots, r_j, \end{aligned} \tag{4}$$

where $\beta_{j\ell 0}$ is the intercept, $\boldsymbol{\beta}_{j\ell i} \in \mathbb{R}^{d_i}$ is the coefficient vector for X_i to predict the ℓ^{th} level of X_j, and $\boldsymbol{\beta}_{j\ell\cdot} = (\beta_{j\ell 0}, \boldsymbol{\beta}_{j\ell 1}, \ldots, \boldsymbol{\beta}_{j\ell p}) \in \mathbb{R}^d$. Note that in the multi-logit model above, $\boldsymbol{\beta}_{j\ell i} = \mathbf{0}$ for all ℓ if $i \notin \Pi_j^{\mathcal{G}}$. We choose to use a symmetric form of the multi-logit model here, as was done in Zhu and Hastie (2004) and Friedman et al. (2010). To make this model identifiable, we impose the following constraints on the intercepts

$$\beta_{j10} = 0, \quad j = 1, \ldots, p. \tag{5}$$

The nonidentifiability of other parameters can be resolved via regularization as demonstrated by Friedman et al. (2010). The particular form of regularization we use leads to the following constraints

$$\sum_{m=1}^{r_j} \boldsymbol{\beta}_{jmi} = \mathbf{0}, \quad \forall\, i, j = 1, \ldots, p. \tag{6}$$

We define $\boldsymbol{\beta}_{j\cdot i} = (\boldsymbol{\beta}_{j1i}, \ldots, \boldsymbol{\beta}_{jr_j i}) \in \mathbb{R}^{d_i r_j}$ to be the vector of coefficients representing the influence of X_i on X_j and $\boldsymbol{\beta}_{j\cdot 0} = (\beta_{j10}, \ldots, \beta_{jr_j 0}) \in \mathbb{R}^{r_j}$ to be the vector of intercepts for predicting X_j. Let $\boldsymbol{\beta} = (\boldsymbol{\beta}_{j\cdot i})_{1 \le j \le p, 0 \le i \le p}$ be the vector of all parameters. Given the structure of $\mathcal{G}$, the number of nonzero elements of $\boldsymbol{\beta}$ is

$$|\text{supp}(\boldsymbol{\beta})| = \sum_{j=1}^{p} \left[(r_j - 1) + r_j \sum_{i \in \Pi_j^{\mathcal{G}}} d_i \right]. \tag{7}$$

If we further assume that $r_i = r$ for all i,

$$|\text{supp}(\boldsymbol{\beta})| = (r-1)\left[r|E| + p\right],$$

which grows linearly in the total number of edges $|E|$. This rate of growth is much slower than that of the product multinomial model (3).

In this paper, we assume the data set $\mathcal{X}$ consists of p blocks, $\mathcal{X}^j$ with size $n_j \times p$, $j = 1, \ldots, p$, where the level of X_j is experimentally fixed in each row of $\mathcal{X}^j$. Let $\mathcal{I}_j$ be the set of row indices of $\mathcal{X}^j$ and $\mathcal{O}_j = \{1, \ldots, n\} \setminus \mathcal{I}_j$ be the index set of rows in which X_j is observational. Denote the size of $\mathcal{O}_j$ by $n_{-j} = |\mathcal{O}_j| = n - n_j$. Under the multi-logit model (4), the log-likelihood function $\ell(\boldsymbol{\beta})$ can be written according to the factorization (2) as

$$\begin{aligned} \ell(\boldsymbol{\beta}) &\propto \sum_{j=1}^{p} \sum_{h \in \mathcal{O}_j} \log \left[p(\mathcal{X}_{hj} | \mathbf{x}_{h,i}, i \in \Pi_j^{\mathcal{G}}) \right] \\ &= \sum_{j=1}^{p} \sum_{h \in \mathcal{O}_j} \left[\sum_{\ell=1}^{r_j} y_{hj\ell} \mathbf{x}_h^T \boldsymbol{\beta}_{j\ell \cdot} - \log \left(\sum_{m=1}^{r_j} \exp(\mathbf{x}_h^T \boldsymbol{\beta}_{jm \cdot}) \right) \right], \end{aligned} \tag{8}$$

where $y_{hj\ell} = I(\mathcal{X}_{hj} = \ell)$ are indicator variables and $\boldsymbol{\beta}_{j\ell k} = \mathbf{0}$ for $k \notin \Pi_j^{\mathcal{G}}$.

Remark 1. Although we have assumed the availability of experimental data, it is easy to see that the log-likelihood (8) applies to observational data as well: If there are no experimental data for X_j, then $\mathcal{O}_j = \{1, \ldots, n\}$ in (8).

2.3 Group Lasso penalty

Estimating the structure of a discrete Bayesian network $\mathcal{G}$ is equivalent to estimating the sparsity pattern of the parameter $\boldsymbol{\beta}$, due to the following equivalence

$$\boldsymbol{\beta}_{j \cdot i} = \mathbf{0} \qquad \Longleftrightarrow \qquad i \notin \Pi_j^{\mathcal{G}}. \tag{9}$$

In order to learn a sparse DAG from data, we estimate $\boldsymbol{\beta}$ via a penalized likelihood approach. It can be seen from (9) that, for discrete Bayesian networks, the set of parents of X_j is given by the set $\{i : \boldsymbol{\beta}_{j \cdot i} \neq \mathbf{0}\}$. The regular Lasso penalty is inappropriate for this purpose since it penalizes each component of $\boldsymbol{\beta}$ separately. We instead wish to penalize the vector $\boldsymbol{\beta}_{j \cdot i} \in \mathbb{R}^{d_i r_j}$ as a whole to obtain a sparse DAG. Yuan and Lin (2006) proposed the group Lasso penalty to select grouped variables (or factors) in linear regression. Subsequently, Kim et al. (2006) extended the

group Lasso to general loss functions and Meier et al. (2008) developed an alternative algorithm for group Lasso penalized logistic regression. Similar to the regular Lasso penalty, the group Lasso penalty suffers certain drawbacks such as inconsistent variable selection under certain circumstances. The adaptive group Lasso was developed to overcome the limitation of the group Lasso (Wang and Leng 2008; Bach 2008; Wei and Huang 2010). Let $\mathcal{G}_{\boldsymbol{\beta}}$ denote the graph induced by $\boldsymbol{\beta}$ so that $\Pi_j^{\mathcal{G}_{\boldsymbol{\beta}}} = \{i : \boldsymbol{\beta}_{j\cdot i} \neq \mathbf{0}\}$ for $j = 1, \ldots, p$. We define the following adaptive group Lasso estimator for a discrete Bayesian network:

$$\hat{\boldsymbol{\beta}}_\lambda \quad = \quad \underset{\boldsymbol{\beta}:\mathcal{G}_{\boldsymbol{\beta}} \text{ is a DAG}}{\arg\min} \left[f_\lambda(\boldsymbol{\beta}) \stackrel{\Delta}{=} -\ell(\boldsymbol{\beta}) + \lambda \sum_{j=1}^{p} \sum_{i=1}^{p} w_{ji} \|\boldsymbol{\beta}_{j\cdot i}\|_2 \right], \tag{10}$$

where $\mathbf{W} = (w_{ji})_{p\times p}$ is a given weight matrix and $\lambda > 0$ is a tuning parameter. See Sections 3.3 and 3.5 for choosing the weights w_{ji} and the parameter λ. Hereafter, we call $\boldsymbol{\beta}_{j\cdot i}$ a (component) group of $\boldsymbol{\beta}$.

3 Algorithm

Parameter estimation for discrete Bayesian networks is computationally demanding because of the nonlinear nature of the multi-logit model (4). We develop in this section a blockwise coordinate descent algorithm to solve (10). Coordinate descent type of algorithms have been proved successful in various settings (Fu 1998; Friedman et al. 2007; Wu and Lange 2008) and their implementations are relatively straightforward.

3.1 Single coordinate descent step

We first consider minimizing $f_\lambda(\boldsymbol{\beta})$ (10) with respect to $\boldsymbol{\beta}_{j\cdot i}$ given the current estimates of all the other parameters. We define

$$\begin{aligned} f_{\lambda,j}(\boldsymbol{\beta}_{j\cdot\cdot}) \quad &= \quad -\sum_{h\in\mathcal{O}_j} \left[\sum_{\ell=1}^{r_j} y_{hj\ell} \mathbf{x}_h^T \boldsymbol{\beta}_{j\ell\cdot} - \log\left(\sum_{m=1}^{r_j} \exp(\mathbf{x}_h^T \boldsymbol{\beta}_{jm\cdot}) \right) \right] + \lambda \sum_{k=1}^{p} w_{jk} \|\boldsymbol{\beta}_{j\cdot k}\|_2 \\ &\stackrel{\Delta}{=} \quad -\ell_j(\boldsymbol{\beta}_{j\cdot\cdot}) + \lambda \sum_{k=1}^{p} w_{jk} \|\boldsymbol{\beta}_{j\cdot k}\|_2, \end{aligned} \tag{11}$$

where $\boldsymbol{\beta}_{j\cdot\cdot} = (\boldsymbol{\beta}_{j\cdot 0}, \boldsymbol{\beta}_{j\cdot 1}, \ldots, \boldsymbol{\beta}_{j\cdot p})$. For simpler notation, we suppress the dependence of $f_{\lambda,j}$ and ℓ_j on $\boldsymbol{\beta}_{j\cdot k}$, $k \neq i$, when considering the problem of minimizing $f_{\lambda,j}(\cdot)$ over $\boldsymbol{\beta}_{j\cdot i}$, and write them as $f_{\lambda,j}(\boldsymbol{\beta}_{j\cdot i})$ and $\ell_j(\boldsymbol{\beta}_{j\cdot i})$, respectively.

Following the approach of Tseng and Yun (2009) and Meier et al. (2008), we form a quadratic approximation to $\ell_j(\boldsymbol{\beta}_{j\cdot i})$ using a second-order Taylor expansion at $\boldsymbol{\beta}_{j\cdot i}^{(t)}$, the current value of $\boldsymbol{\beta}_{j\cdot i}$. Adding the penalty, the quadratic approximation is

$$Q_{\lambda,j}^{(t)}(\boldsymbol{\beta}_{j\cdot i}) = -\left\{(\boldsymbol{\beta}_{j\cdot i} - \boldsymbol{\beta}_{j\cdot i}^{(t)})^T \nabla \ell_j(\boldsymbol{\beta}_{j\cdot i}^{(t)}) + \frac{1}{2}(\boldsymbol{\beta}_{j\cdot i} - \boldsymbol{\beta}_{j\cdot i}^{(t)})^T \mathbf{H}_{ji}^{(t)} (\boldsymbol{\beta}_{j\cdot i} - \boldsymbol{\beta}_{j\cdot i}^{(t)})\right\} + \lambda w_{ji} \|\boldsymbol{\beta}_{j\cdot i}\|_2, \tag{12}$$

up to an additive term that does not depend on $\boldsymbol{\beta}_{j\cdot i}$.

The gradient of the log-likelihood function $\ell_j(\cdot)$ is

$$\nabla \ell_j(\boldsymbol{\beta}_{j\cdot i}^{(t)}) = \sum_{h \in \mathcal{O}_j} \begin{pmatrix} \left(y_{hj1} - p_{j1}^{(t)}(\mathbf{x}_h)\right) \mathbf{x}_{h,i} \\ \vdots \\ \left(y_{hjr_j} - p_{jr_j}^{(t)}(\mathbf{x}_h)\right) \mathbf{x}_{h,i} \end{pmatrix}, \tag{13}$$

where $p_{j\ell}^{(t)}(\mathbf{x}_h)$ defined in (4), $\ell = 1, \ldots, r_j$, are evaluated at the current parameter values. To give a reasonable quadratic approximation, we use a negative definite matrix $\mathbf{H}_{ji}^{(t)} = h_{ji}^{(t)} \mathbf{I}_{d_i r_j}$ in (12) to approximate the Hessian of $\ell_j(\cdot)$, where the scalar $h_{ji}^{(t)} < 0$ and $\mathbf{I}_{d_i r_j}$ is the identity matrix of size $d_i r_j \times d_i r_j$. As suggested by Tseng and Yun (2009) and Meier et al. (2008), we choose

$$h_{ji}^{(t)} = h_{ji}(\boldsymbol{\beta}^{(t)}) \triangleq -\max\{\text{diag}(-\mathbf{H}_{\ell_j}(\boldsymbol{\beta}_{j\cdot i}^{(t)})), b\}, \tag{14}$$

where $\mathbf{H}_{\ell_j}$ is the Hessian of the log-likelihood function $\ell_j(\cdot)$ and b is a small positive number used as a lower bound to help convergence. Note that it is not necessary to recompute $h_{ji}^{(t)}$ every iteration (Meier et al. 2008). See Section 3.4 for more details.

It is not difficult to show the following proposition, which is a direct consequence of the Karush-Kuhn-Tucker (KKT) conditions for minimizing (12).

Proposition 1. *Let* $\mathbf{H}_{ji}^{(t)} = h_{ji}^{(t)} \mathbf{I}_{d_i r_j}$ *for some scalar* $h_{ji}^{(t)} < 0$ *and* $\boldsymbol{d}_{ji}^{(t)} = \nabla \ell_j(\boldsymbol{\beta}_{j\cdot i}^{(t)}) - h_{ji}^{(t)} \boldsymbol{\beta}_{j\cdot i}^{(t)}$.

The minimizer of $Q_{\lambda,j}^{(t)}(\boldsymbol{\beta}_{j\cdot i})$ *in* (12) *is*

$$\bar{\boldsymbol{\beta}}_{j\cdot i}^{(t)} = \begin{cases} \mathbf{0} & \text{if } \|\boldsymbol{d}_{ji}^{(t)}\|_2 \le \lambda w_{ji}, \\ -\dfrac{1}{h_{ji}^{(t)}}\left[\boldsymbol{d}_{ji}^{(t)} - \lambda w_{ji}\dfrac{\boldsymbol{d}_{ji}^{(t)}}{\|\boldsymbol{d}_{ji}^{(t)}\|_2}\right] & \text{otherwise.} \end{cases} \tag{15}$$

In order to achieve sufficient descent, an inexact line search by the Armijo rule is performed when $\bar{\boldsymbol{\beta}}_{j\cdot i}^{(t)} \neq \boldsymbol{\beta}_{j\cdot i}^{(t)}$, following the procedure in Meier et al. (2008). Put

$$\boldsymbol{s}_{ji}^{(t)} = \bar{\boldsymbol{\beta}}_{j\cdot i}^{(t)} - \boldsymbol{\beta}_{j\cdot i}^{(t)}, \tag{16}$$

and let $\Delta^{(t)}$ be the change in $f_{\lambda,j}$ when the log-likelihood is linearized at $\boldsymbol{\beta}_{j\cdot i}^{(t)}$, i.e.,

$$\Delta^{(t)} = -(\boldsymbol{s}_{ji}^{(t)})^T \nabla \ell_j(\boldsymbol{\beta}_{j\cdot i}^{(t)}) + \lambda w_{ji}\|\bar{\boldsymbol{\beta}}_{j\cdot i}^{(t)}\|_2 - \lambda w_{ji}\|\boldsymbol{\beta}_{j\cdot i}^{(t)}\|_2.$$

Pick $0 < \eta < 1$, $0 < \delta < 1$ and $\alpha_0 > 0$, and let $\alpha^{(t)}$ be the largest value in the sequence $\{\alpha_0\eta^k\}_{k\ge 0}$ such that

$$f_{\lambda,j}(\boldsymbol{\beta}_{j\cdot i}^{(t)} + \alpha^{(t)}\boldsymbol{s}_{ji}^{(t)}) - f_{\lambda,j}(\boldsymbol{\beta}_{j\cdot i}^{(t)}) \le \alpha^{(t)}\delta\Delta^{(t)}. \tag{17}$$

Then set

$$\boldsymbol{\beta}_{j\cdot i}^{(t+1)} = \boldsymbol{\beta}_{j\cdot i}^{(t)} + \alpha^{(t)}\boldsymbol{s}_{ji}^{(t)},$$

which completes one iteration for updating $\boldsymbol{\beta}_{j\cdot i}$. In our implementation, $\eta = \delta = 0.9$ and $\alpha_0 = 1$.

Remark 2. It follows from Proposition 1 with $w_{j0} = 0$ that for the unpenalized intercepts,

$$\boldsymbol{\beta}_{j\cdot 0}^{(t+1)} = -\frac{1}{h_{j0}^{(t)}}\boldsymbol{d}_{j0}^{(t)} \tag{18}$$

and the line search is not necessary. In addition, some of the parameters are always constrained to zero, e.g., $\boldsymbol{\beta}_{j\cdot j}$ and β_{j10} for all j. They are not updated.

3.2 Blockwise coordinate descent

Our CD algorithm consists of two layers of iterations. In the outer loop, we cycle through all pairs of nodes to update the active set of edges, including their directions. In the inner loop, we only cycle through the active edge set to update the parameter values.

We first describe the outer loop. Due to acyclicity, we know *a priori* that $\boldsymbol{\beta}_{i\cdot j}$ and $\boldsymbol{\beta}_{j\cdot i}$ cannot simultaneously be nonzero for $i \neq j$. This suggests performing the minimization in blocks, minimizing over $\{\boldsymbol{\beta}_{i\cdot j}, \boldsymbol{\beta}_{j\cdot i}\}$ simultaneously. An immediate consequence of this is a substantial reduction in the number of free parameters. In order to enforce acyclicity, we use a simple heuristic (Fu and Zhou 2013): For each block $\{\boldsymbol{\beta}_{i\cdot j}, \boldsymbol{\beta}_{j\cdot i}\}$, we check to see if adding an edge from $X_i \to X_j$ induces a cycle in the estimated DAG. If so, we set $\boldsymbol{\beta}_{j\cdot i} = \mathbf{0}$ and minimize with respect to $\boldsymbol{\beta}_{i\cdot j}$. Alternatively, if the edge $X_j \to X_i$ induces a cycle, we set $\boldsymbol{\beta}_{i\cdot j} = \mathbf{0}$ and minimize with respect to $\boldsymbol{\beta}_{j\cdot i}$. If neither edge induces a cycle, we minimize over both parameters simultaneously. The cycle check is implemented by a breath-first search algorithm; see the supplementary materials of Fu and Zhou (2013) for more details. We outline below (Algorithm 1) the complete blockwise coordinate descent (CD) algorithm for discrete Bayesian networks. In the algorithm, $\boldsymbol{\beta}_{j\cdot i} \Leftarrow \mathbf{0}$ is used to indicate that $\boldsymbol{\beta}_{j\cdot i}$ must be set to zero due to the acyclicity constraint given the current estimates of the other parameters. Minimization of $f_{\lambda,j}(\cdot)$ with respect to $\boldsymbol{\beta}_{j\cdot i}$ is done with the single CD step in Section 3.1.

Let $\boldsymbol{\beta}^{(t)}$ denote the parameter value after one cycle of the outer loop (after line 19 in Algorithm 1). Denote its active edge set by $E^{(t)} = \{(i,j) : \boldsymbol{\beta}_{j\cdot i}^{(t)} \neq \mathbf{0}\}$. The inner loop solves the following problem:

$$\min_{\boldsymbol{\beta}} f_\lambda(\boldsymbol{\beta}), \text{ subject to } \operatorname{supp}(\boldsymbol{\beta}) \subset E^{(t)}, \tag{19}$$

where f_λ is defined in (10). We use $\boldsymbol{\beta}^{(t)}$ as the initial value and cycle through $\boldsymbol{\beta}_{j\cdot i}$ for $(i,j) \in E^{(t)}$. In particular, the direction of an edge will not be reversed but edges may be deleted if their parameters $\boldsymbol{\beta}_{j\cdot i}$ are updated to zero. See Algorithm 2 for an outline of the inner loop.

By construction, $E^{(t)}$ satisfies the acyclicity constraint and thus the feasible region in (19) is simply a Euclidean space. Since f_λ itself is convex, the CD algorithm for the inner loop has nice convergence properties. In analogy to Proposition 2 in Meier et al. (2008), we arrive at the

Algorithm 1 CD algorithm for estimating discrete Bayesian networks

1: Initialize $\boldsymbol{\beta}$ such that $\mathcal{G}_{\boldsymbol{\beta}}$ is acyclic
2: **for** $i = 1, \ldots, p-1$ **do**
3: **for** $j = i+1, \ldots, p$ **do**
4: **if** $\boldsymbol{\beta}_{j\cdot i} \Leftarrow \mathbf{0}$ **then**
5: $\boldsymbol{\beta}_{i\cdot j} \leftarrow \arg\min_{\boldsymbol{\beta}_{i\cdot j}} f_{\lambda,i}(\cdot), \quad \boldsymbol{\beta}_{j\cdot i} \leftarrow \mathbf{0}$
6: **else if** $\boldsymbol{\beta}_{i\cdot j} \Leftarrow \mathbf{0}$ **then**
7: $\boldsymbol{\beta}_{i\cdot j} \leftarrow \mathbf{0}, \quad \boldsymbol{\beta}_{j\cdot i} \leftarrow \arg\min_{\boldsymbol{\beta}_{j\cdot i}} f_{\lambda,j}(\cdot)$
8: **else**
9: $S_1 \leftarrow \min_{\boldsymbol{\beta}_{i\cdot j}} f_{\lambda,i}(\cdot) + f_{\lambda,j}(\cdot)|_{\boldsymbol{\beta}_{j\cdot i}=\mathbf{0}}$
10: $S_2 \leftarrow f_{\lambda,i}(\cdot)|_{\boldsymbol{\beta}_{i\cdot j}=\mathbf{0}} + \min_{\boldsymbol{\beta}_{j\cdot i}} f_{\lambda,j}(\cdot)$
11: **if** $S_1 \leq S_2$ **then**
12: $\boldsymbol{\beta}_{i\cdot j} \leftarrow \arg\min_{\boldsymbol{\beta}_{i\cdot j}} f_{\lambda,i}(\cdot), \quad \boldsymbol{\beta}_{j\cdot i} \leftarrow \mathbf{0}$
13: **else**
14: $\boldsymbol{\beta}_{i\cdot j} \leftarrow \mathbf{0}, \quad \boldsymbol{\beta}_{j\cdot i} \leftarrow \arg\min_{\boldsymbol{\beta}_{j\cdot i}} f_{\lambda,j}(\cdot)$
15: **end if**
16: **end if**
17: **end for**
18: **end for**
19: Update intercepts $\boldsymbol{\beta}_{j\cdot 0}$ by (18) for $j = 1, \ldots, p$
20: Inner loop given the active set (Algorithm 2)
21: Repeat step 2 to 20 until some stopping criterion is met

Algorithm 2 Inner loop

1: Input $E^{(t)}$ and initialize $\boldsymbol{\beta} = \boldsymbol{\beta}^{(t)}$
2: **for** $(i,j) \in E^{(t)}$ **do**
3: $\bar{\boldsymbol{\beta}}_{j\cdot i} \leftarrow$ Equation (15), $\boldsymbol{s}_{ji} \leftarrow$ Equation (16)
4: $\alpha \leftarrow$ by rule (17)
5: $\boldsymbol{\beta}_{j\cdot i} \leftarrow \boldsymbol{\beta}_{j\cdot i} + \alpha \boldsymbol{s}_{ji}$
6: **end for**
7: Repeat step 2 to step 6 until convergence.

following convergence result.

Proposition 2. *Suppose that the sequence* $\{\boldsymbol{\beta}^{(k)}\}$ *is generated by the inner loop. If the matrix* $\mathbf{H}_{ji}^{(k)}$ *is chosen according to* (14)*, then every limit point of the sequence* $\{\boldsymbol{\beta}^{(k)}\}$ *is a minimizer of problem* (19)*.*

However, since the search space for the outer loop is the full DAG space, which is highly nonconvex, rigorous theory on its convergence is not available. Therefore, a practical stopping criterion is employed. After the convergence of an inner loop, we obtain the current active set. If one more iteration of the outer loop does not change the active set, we then stop Algorithm 1.

On the other hand, we also set a maximum number of iterations for the outer loop. For all the examples we have tested, the stopping criterion is met within five iterations of the outer loop, and our CD algorithm has shown no problem in convergence.

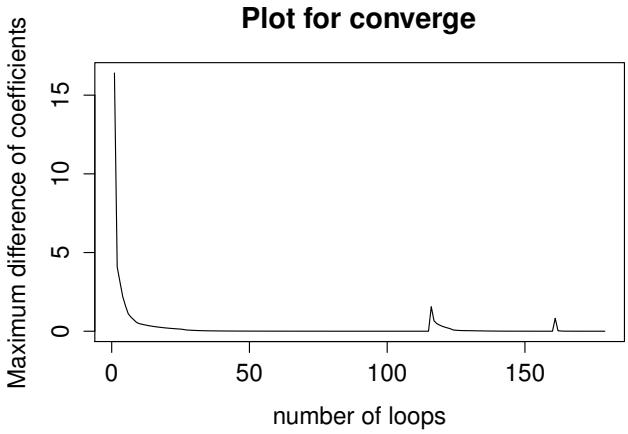

Figure 1: A typical convergence plot

As an illustration, Figure 1 plots $\|\boldsymbol{\beta}^{(t+1)} - \boldsymbol{\beta}^{(t)}\|_\infty$ against t for a dataset generated from a small-world network with $p = 50$, where t indexes iterations in both the inner and the outer loops. In this figure there are three peaks, each appearing after the algorithm finishes an outer loop. The curve between two neighboring peaks demonstrates the convergence of the inner loop. The heights of the three peaks decay very fast, suggesting convergence of the outer loop. Similar patterns are observed for data generated from other types of DAGs.

3.3 Choosing adaptive weights

We choose the weights $w_{ji} = w_{ij} = \min(\|\tilde{\boldsymbol{\beta}}_{j\cdot i}\|_2^{-\gamma}, \|\tilde{\boldsymbol{\beta}}_{i\cdot j}\|_2^{-\gamma})$ for some $\gamma > 0$, where $\tilde{\boldsymbol{\beta}}_{j\cdot i}$ is a consistent estimate of $\boldsymbol{\beta}_{j\cdot i}$. In the generalized linear model setting with a fixed p, one may use the MLEs to define the weights. However, because of the existence of equivalent DAGs, the MLEs may not be consistent in our case. In this work, we construct the initial consistent estimates $\tilde{\boldsymbol{\beta}}_{j\cdot i}$ by minimizing $f_\lambda(\boldsymbol{\beta})$ in (10) with weights $w_{ji} = 1$ for all $i \neq j$, i.e., $\tilde{\boldsymbol{\beta}}_{j\cdot i}$ is simply the group Lasso estimate. Due to the acyclicity constraint, either $\tilde{\boldsymbol{\beta}}_{j\cdot i}$ or $\tilde{\boldsymbol{\beta}}_{i\cdot j}$ or both will be zero. Thus, instead of using $w_{ji} = \|\tilde{\boldsymbol{\beta}}_{j\cdot i}\|_2^{-\gamma}$, we choose $w_{ji} = w_{ij} = \min(\|\tilde{\boldsymbol{\beta}}_{j\cdot i}\|_2^{-\gamma}, \|\tilde{\boldsymbol{\beta}}_{i\cdot j}\|_2^{-\gamma})$

so that the direction of the edge between two nodes will not be pre-determined by the initial estimates. As will be shown in Section 4, our penalized estimator $\hat{\boldsymbol{\beta}}_\lambda$ defined by (10) with this choice of adaptive weights has desired asymptotic properties.

3.4 Solution path

We use Algorithm 1 to compute $\hat{\boldsymbol{\beta}}_\lambda$ (10) over a grid of J values for the tuning parameter, $\lambda_1 > \ldots > \lambda_J > 0$, where at λ_1 every parameter other than the intercepts is estimated as zero. Let $\mathbf{1}_{n_{-j}}$ be the vector of ones of length n_{-j}, $\mathbf{y}_{j\ell} = (y_{hj\ell})_{h\in\mathcal{O}_j} \in \mathbb{R}^{n_{-j}}$, and $\bar{\mathbf{y}}_{j\ell} = \bar{y}_{j\ell}\mathbf{1}_{n_{-j}}$, where $\bar{y}_{j\ell}$ is the mean of $\mathbf{y}_{j\ell}$. Define $\mathbf{X}_{ji}$ to be an $n_{-j} \times d_i$ matrix whose rows are composed of $\mathbf{x}_{h,i}^T$ for $h \in \mathcal{O}_j$. It follows from the KKT conditions for (10) that

$$\lambda_1 = \max_{1\le i,j\le p} \frac{1}{w_{ji}} \|\nabla \ell_j(\boldsymbol{\beta}_{j\cdot i})|_{\boldsymbol{\beta}_{j\cdot i}=\mathbf{0}}\|_2 = \max_{1\le i,j\le p} \frac{1}{w_{ji}} \left\| \begin{pmatrix} \mathbf{X}_{ji}^T(\mathbf{y}_{j1} - \bar{\mathbf{y}}_{j1}) \\ \vdots \\ \mathbf{X}_{ji}^T(\mathbf{y}_{jr_j} - \bar{\mathbf{y}}_{jr_j}) \end{pmatrix} \right\|_2, \tag{20}$$

where $\boldsymbol{\beta}_{j\cdot 0}$ are set to the MLEs of the intercepts assuming all $\boldsymbol{\beta}_{j\cdot i}$, $i = 1, \ldots, p$, are zero.

The solution $\hat{\boldsymbol{\beta}}_{\lambda_m}$ is used as a warm start for estimating $\hat{\boldsymbol{\beta}}_{\lambda_{m+1}}$, $m = 1, \ldots, J-1$. To save computational time, we set $h_{ji}^{(t)} = h_{ji}(\hat{\boldsymbol{\beta}}_{\lambda_m})$ (14) in the CD algorithm for $\hat{\boldsymbol{\beta}}_{\lambda_{m+1}}$, instead of updating $h_{ji}^{(t)}$ every iteration.

3.5 Tuning parameter selection

Traditional model selection criteria such as cross-validation and BIC do not work well for the purpose of estimating DAGs from data, because the DAG with the smallest prediction error often has much more edges than the true DAG. In order to select a suitable tuning parameter λ_m, we use an empirical model selection criterion proposed by Fu and Zhou (2013). Let $\hat{\mathcal{G}}_{\lambda_m}$ be the DAG induced by $\hat{\boldsymbol{\beta}}_{\lambda_m}$ and e_{λ_m} be the number of directed edges in $\hat{\mathcal{G}}_{\lambda_m}$. We reestimate $\boldsymbol{\beta}$ by the maximizer $\boldsymbol{\beta}_{\lambda_m}^\dagger$ of the log-likelihood $\ell(\boldsymbol{\beta})$ (8) given $\mathcal{G} = \hat{\mathcal{G}}_{\lambda_m}$, in which regressing of X_j on $\Pi_j^{\mathcal{G}}$ is done using the R package *mlogit* (Croissant 2011). We define the difference ratio between two estimated DAGs $\hat{\mathcal{G}}_{\lambda_m}$ and $\hat{\mathcal{G}}_{\lambda_{m+1}}$ by $dr_{(m,m+1)} = \Delta\ell_{(m,m+1)}/\Delta e_{(m,m+1)}$, where

$\Delta\ell_{(m,m+1)} = \ell(\boldsymbol{\beta}^{\dagger}_{\lambda_{m+1}}) - \ell(\boldsymbol{\beta}^{\dagger}_{\lambda_m})$ and $\Delta e_{(m,m+1)} = e_{\lambda_{m+1}} - e_{\lambda_m}$, provided that $\Delta e_{(m,m+1)} \geq 1$. Otherwise, we set $dr_{(m,m+1)} = dr_{(m-1,m+1)}$. The selected tuning parameter is indexed by

$$m^* = \sup\left\{2 \leq m \leq J : dr_{(m-1,m)} \geq \alpha \cdot \max\{dr_{(1,2)}, \ldots, dr_{(J-1,J)}\}\right\}. \tag{21}$$

The typical value of the thresholding parameter α is 0.1. According to this criterion, an increase in model complexity, measured by the number of predicted edges, is accepted only if there is a substantial increase in the log-likelihood.

4 Asymptotic Properties

In this section we establish asymptotic theory for the DAG estimator $\hat{\boldsymbol{\beta}}_\lambda$ assuming that p is fixed and $n \to \infty$. By rearranging and relabeling individual components, we rewrite $\boldsymbol{\beta}$ as $\boldsymbol{\phi} = (\boldsymbol{\phi}_{(1)}, \boldsymbol{\phi}_{(2)})$, where $\boldsymbol{\phi}_{(1)} = (\boldsymbol{\beta}_{1\cdot 1}, \ldots, \boldsymbol{\beta}_{1\cdot p}, \ldots\ldots, \boldsymbol{\beta}_{p\cdot 1}, \ldots, \boldsymbol{\beta}_{p\cdot p})$ is the parameter vector of interest and $\boldsymbol{\phi}_{(2)} = (\boldsymbol{\beta}_{1\cdot 0}, \ldots, \boldsymbol{\beta}_{p\cdot 0})$ denotes the vector of intercepts. Hereafter, we denote by ϕ_j the j^{th} group of $\boldsymbol{\phi}$, such that $\phi_1 = \boldsymbol{\beta}_{1\cdot 1}$, $\phi_2 = \boldsymbol{\beta}_{1\cdot 2}$, $\ldots$, $\phi_{p^2} = \boldsymbol{\beta}_{p\cdot p}$, and so on. We concatenate the rows of the weight matrix $\mathbf{W}$ into a vector $\mathcal{T} = (\tau_j)_{1:p^2} = (w_{11}, \ldots, w_{1p}, \ldots\ldots, w_{p1}, \ldots, w_{pp})$. We say $\boldsymbol{\phi}$ is acyclic if the graph $\mathcal{G}_{\boldsymbol{\phi}}$ induced by $\boldsymbol{\phi}$ (or the corresponding $\boldsymbol{\beta}$) is acyclic. Let $\Omega = \{\boldsymbol{\phi} : \boldsymbol{\phi}$ is acyclic and satisfies (5) and (6)$\}$ denote the parameter space, $\boldsymbol{\phi}^* \in \Omega$ be the true parameter, and $\mathcal{G}_{\boldsymbol{\phi}^*}$ be the true DAG. Let $\mathcal{A} = \{j : \phi_j^* = \mathbf{0}, 1 \leq j \leq p^2\}$ and $\mathcal{B} = \{j : \phi_j^* \neq \mathbf{0}, 1 \leq j \leq p^2\}$ denote the groupwise sparsity pattern of $\boldsymbol{\phi}^*$.

Define $\boldsymbol{\phi}_{[k]}$ ($k \in \{1, \ldots, p\}$) to be the parameter vector obtained from $\boldsymbol{\phi}$ by setting $\boldsymbol{\beta}_{k\cdot i} = \mathbf{0}$ for $i = 1, \ldots, p$. In other words, the DAG $\mathcal{G}_{\boldsymbol{\phi}_{[k]}}$ is obtained by deleting all edges pointing to the k^{th} node in $\mathcal{G}_{\boldsymbol{\phi}}$. As demonstrated by (2), we can model interventional data in the k^{th} block of the data matrix $\mathcal{X}^k$ as *i.i.d.* observations from a joint distribution factorized according to $\mathcal{G}_{\boldsymbol{\phi}_{[k]}}$. Denote the corresponding probability mass function by $p(\mathbf{x}|\boldsymbol{\phi}_{[k]})$, where $\mathbf{x} = (x_1, \ldots, x_p)$ and $x_j \in \{1, \ldots, r_j\}$ for $j = 1, \ldots, p$. The log-likelihood of the full data matrix $\mathcal{X}$ is

$$L(\boldsymbol{\phi}) = \sum_{k=1}^{p} L_k(\boldsymbol{\phi}_{[k]}) = \sum_{k=1}^{p} \log p(\mathcal{X}^k \mid \boldsymbol{\phi}_{[k]}), \tag{22}$$

where $\log p(\mathcal{X}^k|\boldsymbol{\phi}_{[k]}) = \sum_{h\in\mathcal{I}_k} \log(p(\mathcal{X}_{h\cdot}|\boldsymbol{\phi}_{[k]}))$ and $\mathcal{X}_{h\cdot} = (\mathcal{X}_{h1}, \ldots, \mathcal{X}_{hp})$. Then the penalized log-likelihood function with the adaptive group Lasso penalty is

$$R(\boldsymbol{\phi}) = L(\boldsymbol{\phi}) - \lambda_n \sum_{j=1}^{p^2} \tau_j \|\phi_j\|_2 = \sum_{k=1}^{p} L_k(\boldsymbol{\phi}_{[k]}) - \lambda_n \sum_{j=1}^{p^2} \tau_j \|\phi_j\|_2, \tag{23}$$

where the component group ϕ_j ($j = 1, \ldots, p^2$) represents the influence of one variable on another. A penalized likelihood estimator $\hat{\boldsymbol{\phi}}$ is obtained by maximizing $R(\boldsymbol{\phi})$ in the parameter space Ω.

Though interventional data help to distinguish equivalent DAGs, the following notion of natural parameters is needed to completely establish identifiability of DAGs for the case where each variable has interventional data. We say that X_i is an ancestor of X_j in a DAG $\mathcal{G}$ if there exists at least one directed path from X_i to X_j. Denote the set of ancestors of X_j by $\text{an}(X_j)$.

Definition 1 (Natural parameters)**.** We say that $\boldsymbol{\phi} \in \Omega$ is natural if $X_i \in \text{an}(X_j)$ in $\mathcal{G}_{\boldsymbol{\phi}}$ implies that X_j is not independent of X_i according to the joint distribution given by the DAG $\boldsymbol{\phi}_{[i]}$ for all $i, j = 1, \ldots, p$.

For a causal DAG, a natural parameter implies that the effects along multiple causal paths connecting the same pair of nodes do not cancel, which is a reasonable assumption for many real-world problems. To establish asymptotic properties of our penalized likelihood estimator for fixed p, we make the following assumptions:

(A1) The true parameter $\boldsymbol{\phi}^*$ is natural and a fixed interior point of Ω.

(A2) The conditional distribution $[X_j|\Pi_j^{\mathcal{G}_{\boldsymbol{\phi}}}]$, $\boldsymbol{\phi} \in \Omega$, is specified by the multi-logit model (4) with constraints (5) and (6) for $j = 1, \ldots, p$.

Recall that the k^{th} block of our data, $\mathcal{X}^k$, can be regarded as an *i.i.d.* sample of size n_k from the distribution $p(\mathbf{x}|\boldsymbol{\phi}^*_{[k]})$ for $k = 1, \ldots, p$. Proofs of the following theorems are relegated to the Appendix.

Theorem 1. *Assume (A1) and (A2). If* $p(\mathbf{x}|\boldsymbol{\phi}_{[k]}) = p(\mathbf{x}|\boldsymbol{\phi}^*_{[k]})$ *for all possible* $\mathbf{x}$ *and all* $k = 1, \ldots, p$*, then* $\boldsymbol{\phi} = \boldsymbol{\phi}^*$*. Furthermore, if* $n_k/n \to \alpha_k > 0$*, then for any* $\boldsymbol{\phi} \neq \boldsymbol{\phi}^*$

$$P(L(\boldsymbol{\phi}^*) > L(\boldsymbol{\phi})) \to 1 \quad \text{as } n \to \infty. \tag{24}$$

Theorem 2. *Assume (A1), (A2), and the weights* $\tau_j = O_p(1)$ *in* (23) *for all* j. *If* $\lambda_n/\sqrt{n} \to 0$ *and* $n_k/n \to \alpha_k > 0$ *for* $k = 1, \ldots, p$, *then there exists a global maximizer* $\hat{\boldsymbol{\phi}}$ *of* $R(\boldsymbol{\phi})$ *such that* $\|\hat{\boldsymbol{\phi}} - \boldsymbol{\phi}^*\|_2 = O_p(n^{-1/2})$.

Recall that the adaptive weight we use is $w_{ji} = w_{ij} = \min(\|\tilde{\boldsymbol{\beta}}_{j\cdot i}\|_2^{-\gamma}, \|\tilde{\boldsymbol{\beta}}_{i\cdot j}\|_2^{-\gamma})$. For $m \in \{1, \ldots, p^2\}$, define an index $r(m)$ so that

$$\phi_m = \boldsymbol{\beta}_{j\cdot i} \text{ and } \phi_{r(m)} = \boldsymbol{\beta}_{i\cdot j} \tag{25}$$

for some $i, j \in \{1, \ldots, p\}$. That is, ϕ_m and $\phi_{r(m)}$ index the parameters for a pair of reversed edges (if $i \neq j$).

Theorem 3. *Assume (A1) and (A2). Let the weights* $\tau_j = \min(\|\tilde{\phi}_j\|_2^{-\gamma}, \|\tilde{\phi}_{r(j)}\|_2^{-\gamma})$ *in* (23) *for all* j, *where* $\gamma > 0$ *and* $\tilde{\phi}_j$ *is* $\sqrt{n}$*-consistent for* ϕ_j^*. *If* $\lambda_n/\sqrt{n} \to 0$, $\lambda_n n^{(\gamma-1)/2} \to \infty$ *and* $n_k/n \to \alpha_k > 0$ *for* $k = 1, \ldots, p$, *then with probability tending to one, there exists a global maximizer* $\hat{\boldsymbol{\phi}}$ *of* $R(\boldsymbol{\phi})$ *such that* $\|\hat{\boldsymbol{\phi}} - \boldsymbol{\phi}^*\|_2 = O_p(n^{-1/2})$ *and* $\hat{\boldsymbol{\phi}}_{\mathcal{A}} = (\hat{\phi}_j)_{j\in\mathcal{A}} = \mathbf{0}$.

Theorem 1 confirms that the causal DAG model is identifiable with interventional data assuming a natural parameter. Theorem 2 implies that there is a $\sqrt{n}$-consistent global maximizer of $R(\boldsymbol{\phi})$ with unweighted group Lasso penalty ($\tau_j = 1$, $\forall j$), which in turn can be used to construct adaptive weights to achieve model selection consistency (Theorem 3).

Remark 3. It is interesting to generalize the above asymptotic results for fixed p to the case where $p = p_n$ grows with the sample size n, say, by developing nonasymptotic bounds on the ℓ_2 estimation error $\|\hat{\boldsymbol{\phi}} - \boldsymbol{\phi}^*\|_2$. However, in order to estimate the causal network consistently, sufficient intervention data are needed for each node, i.e., n_k must approach infinity, and thus $p/n \to 0$ as $n \to \infty$. This limits us to the low-dimensional setting with $p < n$. Suppose we have a large network with $p \gg n$. One may first apply some regularization method on observational data to screen out independent nodes and to partition the network into small subgraphs that are disconnected to each other. Then for each small subgraph, we can afford to generate enough interventional data for each node and apply the method in this paper to infer

the causal structure. Our asymptotic theory provides useful guidance for the analysis in the second step.

For purely observational data, the theory becomes more complicated due to the existence of equivalent DAGs and parameterizations. In general, there exist multiple DAGs that code the same true distribution for the data. Let $\boldsymbol{\phi}^*$ be such a DAG. By a similar proof, one can establish that there is a *local* maximizer $\hat{\boldsymbol{\phi}}$ of the penalized likelihood $R(\boldsymbol{\phi})$ such that $\|\hat{\boldsymbol{\phi}} - \boldsymbol{\phi}^*\|_2 = O_p(n^{-1/2})$ under comparable conditions of Theorem 2. Moreover, if adaptive weights are used as in Theorem 3, the local maximizer $\hat{\boldsymbol{\phi}}$ is model selection consistent for $\boldsymbol{\phi}^*$. It is left as future work to establish the consistency of a global maximizer for high-dimensional observational data and to show that the maximizer has a comparable sparsity as the sparsest true DAG. For the second problem, concave penalties may be needed according to our preliminary analysis on Gaussian Bayesian networks (manuscript in preparation).

5 Simulation Study

We evaluate the CD algorithm on simulated data sets. As stated in Remark 1, the log-likelihood (8) applies to observational data as well: If there are no experimental data for X_j, then $\mathcal{O}_j = \{1, \ldots, n\}$ in (8). Therefore, we applied the CD algorithm to both interventional data and observational data. In order to assess the accuracy and efficiency of the CD algorithm, we compare our algorithm with a few competing methods. For interventional data, we compare our CD algorithm with a Monte Carlo method, the equi-energy sampler (EE sampler) (Kou et al. 2006; Zhou 2011) and a constraint-based method, the PC algorithm (Kalisch and Bühlmann 2007). For observational data we compare the CD algorithm with a score-based method, the hill-climbing (HC) algorithm (Russell and Norvig 1995), and also the PC algorithm. We are interested in comparing both the accuracy and the speed of each algorithm.

Details about data generation and parameter choices will be discussed in Section 5.1. In Section 5.2, we compare DAGs and their skeletons estimated from interventional data. Section 5.3, on the other hand, presents results on high-dimensional observational data. The comparison of running times is provided in Section 5.4.

5.1 Experimental setup

Four types of networks are used to compare the methods: the bipartite graph, the polytree, the scale-free network, and the small-world network. In each setting, we consider the combination of three main parameters: (n, p, s_0), where n is the sample size, p is the number of nodes, and s_0 is the number of true edges.

Bipartite graphs, scale-free networks, and small-world networks were generated using the R package *igraph* (Csárdi and Nepusz 2006). The bipartite graphs were generated by the Erdős-Rényi model (Renyi and Erdos 1959). Each bipartite graph in our datasets has $0.2p$ top nodes, $0.8p$ bottom nodes, and $s_0 = p$ directed edges from the top to the bottom nodes. The structure of a scale-free network was generated using the Barabási-Albert model (Barabási and Albert 1999). These networks have $s_0 = p - 1$ directed edges. The small-world networks were generated using the Watts-Strogatz model (Watts and Strogatz 1998). A graph initially generated by the model was undirected. To convert it to a DAG, edge directions were chosen according to a randomly generated topological sort (Appendix). In this way, a small-world network has $s_0 = 2p$ directed edges. To obtain a polytree, we first generated five binary trees, each of $0.2p$ nodes. Then, the five binary trees were randomly connected via directed edges to form a polytree of $s_0 = p - 1$ directed edges.

In all the simulation studies, each variable is assumed to be binary, i.e., $r_j = 2$ for all j. In this case, each group of parameters $\boldsymbol{\beta}_{j\cdot i} = (\beta_{j1i}, \beta_{j2i}) \in \mathbb{R}^2$. If $\Pi_j = \varnothing$, X_j will be sampled from its two levels with equal probability. Otherwise, the parameters $\boldsymbol{\beta}_{j\cdot 0}$ and $\boldsymbol{\beta}_{j\cdot i}$, $i \in \Pi_j$, are chosen such that

$$p_{j\ell}(\mathbf{x}_h) = \frac{\exp(2\sum_{i\in\Pi_j} y_{hi\ell})}{\exp(2\sum_{i\in\Pi_j} y_{hi1}) + \exp(2\sum_{i\in\Pi_j} y_{hi2})}$$

for $\ell = 1, 2$, where $y_{hi\ell} = I(\mathcal{X}_{hi} = \ell)$. The value of a variable under intervention is randomly fixed to one of its levels regardless of its parents, while other variables are generated according to the multi-logit model (4) given a DAG structure.

We first ran our CD algorithm with weights $w_{ji} = 1$, $\forall i, j$, to obtain an initial estimate $\tilde{\boldsymbol{\beta}}$, which was then used to calculate the adaptive Lasso weights as in Section 3.3 with $\gamma = 1$. For each dataset, we generated a sequence of 30 values of λ, starting from λ_1 (20) and ending at

$0.1\lambda_1$. Consequently, 30 DAGs were estimated, and model selection was done using criterion (21) with $\alpha = 0.1$.

The EE sampler was used for comparisons on interventional data. Its implementation for DAG estimation from interventional data was done as in Zhou (2011). We will call it the EE-DAG sampler hereafter. For each run, we simulated 10 chains with a 0.1 chance of equi-energy jumps. To obtain an estimated DAG, we thresholded the average graph from the target chain (the 10^{th} chain). More precisely, a directed edge was predicted if its posterior inclusion probability was greater than 0.5.

The HC algorithm is a standard greedy method and also a score-based method. For this algorithm, we used the R package `bnlearn` (Scutari 2010). This method is designed especially for observational data and thus was compared with our method only on observational data. Since there are no tuning parameters that strongly influence the accuracy of the final estimates, we simply used its default settings.

The PC algorithm is a popular constraint-based algorithm for learning Bayesian networks, with an efficient implementation in the R package `pcalg` (Kalisch et al. 2012). In the comparisons, we applied it to both interventional and observational data. Since the PC algorithm cannot handle interventional data directly and cannot generate a DAG in general, we consider a two-step approach, called the PC-based method, to produce results favorable for the PC algorithm for interventional data. The output from the PC algorithm is a completed partially directed acyclic graph (CPDAG), which contains both directed and undirected edges. We first apply the PC algorithm to estimate a CPDAG from data. Then one may try to estimate the direction of an undirected edge using intervention data and produce a DAG. In the comparisons in Section 5.2, however, we simply count an undirected edge between two nodes in a CPDAG as an expected edge, provided that there is a corresponding directed edge between the two nodes in the true DAG. Thus, the result reported for the PC-based method is an upper bound performance one can obtain for estimating a causal DAG. For observational data (Section 5.3), we only consider the skeleton from a CPDAG for comparison. The tuning parameter for the PC algorithm is the significance level for conditional independence tests, which was chosen as $\alpha = 0.05$ for all data.

5.2 Results for interventional data

For each type of network we generated two graphs, with $p = 50$ and $p = 100$ nodes, respectively. Recall that the sample size $n = pn_j$ for interventional data, where n_j is the size of the j^{th} block, $\mathcal{X}^j$ $(j = 1, \ldots, p)$. We chose $n_j \in \{1, 5\}$ to test the performance of the algorithms given different amount of intervention. Particularly, when $n_j = 1$ we have $n = p$, which lies on the boundary between low- and high-dimensional settings. In combination, our choices of the data size are $(n, p) \in \{(50, 50), (250, 50), (100, 100), (500, 100)\}$. The number of true edges s_0 for each graph is given in Tables 1 and 2. For each combination of (n, p, s_0), we generated 20 data sets.

Table 1: Average results for DAGs learned from simulated interventional data ($n = p$)

Graph	(n, p, s_0)	Method	P	E	R	FP	TPR	FDR
Bipartite	(50, 50, 50)	CD	39.9	21.1	9.0	9.8	0.422(0.074)	0.470(0.080)
		EE	54.5	27.6	5.0	21.8	0.553(0.080)	0.493(0.102)
		CD*	27.4	16.4	7.8	3.2	0.328(0.056)	0.402(0.106)
		PC	27.4	2.0(11.1)	7.2	7.0	0.262(0.043)	0.521(0.098)
	(100, 100, 100)	CD	82.2	53.1	22.0	7.1	0.531(0.039)	0.354(0.048)
		CD*	65.4	43.2	18.1	4.2	0.432(0.036)	0.341(0.054)
		PC	65.2	9.0(16.6)	27.5	12.0	0.256(0.040)	0.607(0.062)
Polytree	(50, 50, 49)	CD	50.3	24.8	17.0	8.5	0.507(0.085)	0.506(0.076)
		EE	60.8	33.6	7.2	20.0	0.686(0.077)	0.448(0.090)
		CD*	35.6	20.1	13.3	2.2	0.410(0.067)	0.436(0.082)
		PC	35.2	3.8(22.8)	6.4	2.2	0.542(0.065)	0.246(0.078)
	(100, 100, 99)	CD	105.4	58.6	35.4	11.4	0.591(0.067)	0.444(0.076)
		CD*	91.4	54.1	32.2	5.0	0.546(0.063)	0.408(0.073)
		PC	91.0	17.6(62.4)	8.1	2.8	0.808(0.040)	0.120(0.044)
Scale-free	(50, 50, 49)	CD	55.4	26.3	10.0	19.1	0.537(0.086)	0.525(0.085)
		EE	60.8	30.5	4.0	26.2	0.621(0.092)	0.499(0.077)
		CD*	29.2	17.3	7.6	4.2	0.353(0.066)	0.406(0.090)
		PC	28.0	3.2(15.0)	0.6	9.1	0.372(0.040)	0.348(0.089)
	(100, 100, 99)	CD	107.0	64.4	26.2	16.4	0.650(0.063)	0.398(0.069)
		CD*	70.0	47.0	20.2	2.8	0.475(0.055)	0.328(0.069)
		PC	69.6	9.7(45.9)	3.4	10.6	0.562(0.032)	0.201(0.055)
Small-world	(50, 50, 100)	CD	52.3	24.3	23.2	4.8	0.242(0.038)	0.536(0.069)
		EE	59.3	31.0	14.8	13.6	0.310(0.038)	0.477(0.060)
		CD*	34.7	16.2	16.6	2.0	0.162(0.036)	0.534(0.097)
		PC	34.7	7.4(20.0)	5.4	2.0	0.274(0.024)	0.210(0.069)
	(100, 100, 200)	CD	107.0	51.8	50.4	4.8	0.259(0.020)	0.516(0.037)
		CD*	91.3	44.5	43.7	3.1	0.222(0.015)	0.512(0.030)
		PC	91.4	29.4(42.4)	16.4	3.2	0.359(0.026)	0.214(0.051)

Note: Numbers in parentheses in the columns TPR and FDR are the standard deviations across 20 data sets. The number in parentheses in column E for the PC-based method reports the number of predicted undirected edges in the true skeletion.

Table 2: Average results for DAGs learned from simulated interventional data ($n = 5p$)

Graph	(n, p, s_0)	method	P	E	R	FP	TPR	FDR
Bipartite	$(250, 50, 50)$	CD	48.8	35.7	10.4	2.7	0.714(0.070)	0.268(0.073)
		EE	51.3	48.4	1.4	1.6	0.967(0.024)	0.058(0.032)
		CD*	35.8	25.8	8.9	1.2	0.515(0.067)	0.281(0.073)
		PC	35.4	9.8(11.9)	9.6	4.2	0.433(0.041)	0.388(0.074)
	$(500, 100, 100)$	CD	99.6	75.0	22.2	2.4	0.750(0.040)	0.247(0.040)
		CD*	92.8	70.2	21.0	1.6	0.702(0.039)	0.243(0.040)
		PC	92.4	39.2(25.4)	16.4	11.3	0.646(0.042)	0.300(0.052)
Polytree	$(250, 50, 49)$	CD	52.6	30.0	17.5	5.1	0.612(0.102)	0.429(0.114)
		EE	50.3	48.4	0.6	1.3	0.988(0.019)	0.038(0.028)
		CD*	50.0	29.6	17.0	3.3	0.605(0.101)	0.406(0.101)
		PC	49.8	17.4(29.0)	2.3	1.1	0.948(0.033)	0.068(0.046)
	$(500, 100, 99)$	CD	109.3	60.2	37.6	11.5	0.608(0.042)	0.450(0.039)
		CD*	101.7	59.3	36.6	5.8	0.599(0.042)	0.417(0.037)
		PC	101.5	43.8(54.3)	1.0	2.5	0.990(0.014)	0.034(0.021)
Scale-free	$(250, 50, 49)$	CD	52.9	38.4	10.2	4.3	0.783(0.086)	0.274(0.098)
		EE	50.2	48.3	0.7	1.2	0.986(0.013)	0.037(0.025)
		CD*	35.4	27.2	8.0	0.2	0.555(0.091)	0.232(0.105)
		PC	35.1	3.3(25.6)	0.3	5.9	0.590(0.023)	0.177(0.050)
	$(500, 100, 99)$	CD	107.3	74.3	24.6	8.4	0.750(0.045)	0.308(0.061)
		CD*	85.2	63.0	22.0	0.2	0.637(0.047)	0.260(0.044)
		PC	85.0	12.1(60.4)	1.5	11.0	0.732(0.016)	0.147(0.043)
Small-world	$(250, 50, 100)$	CD	59.0	31.5	26.1	1.4	0.315(0.041)	0.466(0.068)
		EE	80.0	64.7	10.2	5.1	0.647(0.072)	0.192(0.067)
		CD*	58.2	31.2	25.8	1.3	0.312(0.044)	0.464(0.070)
		PC	59.0	24.6(14.6)	19.5	0.3	0.392(0.047)	0.335(0.083)
	$(500, 100, 200)$	CD	119.6	63.8	53.2	2.5	0.319(0.037)	0.466(0.052)
		CD*	119.9	64.1	53.3	2.5	0.320(0.038)	0.465(0.053)
		PC	138.6	74.0(23.9)	39.8	0.9	0.489(0.029)	0.293(0.040)

In this subsection, we compare the DAGs estimated by three algorithms, the CD algorithm, the EE-DAG sampler, and the PC-based method. For an estimated DAG, we distinguish between expected edges, which are estimated edges in the true skeleton with the correct direction, and reversed edges, which are in the true skeleton but with a reversed direction. Let P, E, R, and FP denote, respectively, the number of predicted edges, the number of expected edges, the number of reversed edges, and the number of false positive edges (excluding the reversed ones) in an estimated DAG, and let T be the number of edges in the true graph. The accuracy of DAG estimation is measured by the true positive rate (TPR) and the false discovery rate (FDR), defined as TPR = E/T and FDR = (R + FP)/P. Results reported in Table 1 ($n = p$) and Table 2 ($n = 5p$) are averages over 20 data sets for each of the eight simulated networks. To

ease comparison, from the solution path generated by the CD algorithm we selected the DAG with the closest match to the number of edges predicted by the PC algorithm, and report this result in the tables (CD*) as well. From the two tables, it is obvious that the larger the n_j is, the better the results are, especially for the EE-DAG sampler and the PC algorithm.

We were only able to test the EE-DAG sampler on small graphs with $p = 50$ because its computing time was too long for graphs with $p = 100$. For the cases of $n = p = 50$ (Table 1), the EE-DAG sampler predicted much more false positive edges (FP) but fewer reversed edges (R) than the CD algorithm. Its FDR (which accounts for both FP and R) turns out to be comparable to our CD algorithm, but the EE-DAG sampler always predicted the highest number of expected edges. For the cases of $n = 5p$ in Table 2, the EE-DAG sampler has an outstanding performance, with the highest number of expected edges and also the least number of reversed and false positive edges. These observations are largely in agreement with our expectation. The EE-DAG sampler searches for DAGs by sampling from a posterior distribution under the product multinomial model with a conjugate Dirichlet prior, which is close to ℓ_0 regularization when n is large. Thus, it is expected to have good performance when p is small. However, the search is combinatorial in nature, which makes it impractical for even moderately large networks (such as the graphs with $p = 100$ here). On the contrary, our CD algorithm showed no problem in estimating graphs with $p = 100$ nodes and can obtain comparable or even better results than the cases of $p = 50$.

The number of predicted edges by our model selection criterion is quite close to the number of true edges s_0 for most of the graphs. In general, criterion (21) tends to pick DAGs with more edges than the ones estimated by the PC algorithm. We focus on the comparison between the two methods after the edge numbers have been matched, i.e., between CD* and PC in the tables. For almost all graphs, our method estimated fewer number of false positive edges (FP) and more true directed edges than the PC algorithm. These advantages are more substantial when the sample size is relatively small ($n = p$, Table 1). Our CD algorithm predicted more reversed edges for some graphs, such as the small-world networks and polytrees. However, one also sees a large number of undirected edges from the PC algorithm for these cases. When calculating the TPRs and FDRs in the tables, these undirected edges were counted as expected

Table 3: Average results for skeleton estimation from simulated interventional data

		CD Algorithm			PC Algorithm		
Graph	(n, p, s_0)	P	TPR	FDR	P	TPR	FDR
Bipartite	$(50, 50, 50)$	27.4	0.484(0.040)	0.118(0.064)	27.4	0.407(0.037)	0.256(0.068)
	$(100, 100, 100)$	65.4	0.612(0.038)	0.064(0.039)	65.2	0.531(0.017)	0.185(0.029)
Polytree	$(50, 50, 49)$	35.6	0.682(0.065)	0.063(0.041)	35.2	0.673(0.046)	0.062(0.032)
	$(100, 100, 99)$	91.4	0.872(0.030)	0.055(0.031)	91.0	0.890(0.022)	0.031(0.017)
Scale-free	$(50, 50, 49)$	29.2	0.508(0.060)	0.146(0.055)	28.0	0.386(0.036)	0.325(0.084)
	$(100, 100, 99)$	70.0	0.679(0.032)	0.040(0.031)	69.6	0.596(0.017)	0.152(0.037)
Small-world	$(50, 50, 100)$	34.7	0.327(0.020)	0.058(0.038)	34.7	0.328(0.017)	0.056(0.034)
	$(100, 100, 200)$	91.3	0.441(0.018)	0.034(0.014)	91.4	0.441(0.018)	0.035(0.018)
Bipartite	$(250, 50, 50)$	35.8	0.693(0.057)	0.032(0.034)	35.4	0.624(0.026)	0.119(0.036)
	$(500, 100, 100)$	92.8	0.912(0.031)	0.017(0.013)	92.4	0.811(0.020)	0.122(0.018)
Polytree	$(250, 50, 49)$	50.0	0.952(0.024)	0.066(0.023)	49.8	0.995(0.009)	0.022(0.019)
	$(500, 100, 99)$	101.7	0.969(0.012)	0.057(0.014)	101.5	1.000(0.000)	0.025(0.014)
Scale-free	$(250, 50, 49)$	35.4	0.717(0.050)	0.007(0.015)	35.1	0.596(0.022)	0.168(0.046)
	$(500, 100, 99)$	85.2	0.858(0.044)	0.002(0.005)	85.0	0.747(0.011)	0.130(0.036)
Small-world	$(250, 50, 100)$	58.2	0.570(0.024)	0.022(0.022)	59.0	0.588(0.028)	0.005(0.010)
	$(500, 100, 200)$	119.9	0.587(0.020)	0.021(0.010)	138.6	0.688(0.015)	0.006(0.008)

ones, which clearly favors the PC-based method.

To complement the above results, we also compare the two methods in terms of estimating the skeleton of a DAG. In this case, we do not distinguish expected edges from reversed ones and correspondingly, the TPR and FDR for skeleton estimation are defined respectively as TPR = E/T and FDR = FP/P, where E now denotes the number of estimated edges that are in the true skeleton (including reversed edges). Note that an estimated CPDAG from the PC algorithm is sufficient for determining the skeleton. Summary results for skeleton estimation by the two methods are shown in Table 3, in which the top and the bottom panels report the results for $n = p$ and $n = 5p$, respectively. To ease comparison, we have again matched the number of predicted edges of our CD algorithm with the PC algorithm. For bipartite and scale-free graphs, our CD algorithm show much better performance than the PC algorithm, with higher TPRs and lower FDRs. For the other two types of graphs, the two methods have almost identical performance for the datasets with $n = p$, while the PC algorithm outperforms our method slightly for $n = 5p$. Overall, our CD algorithm has more advantages over the PC algorithm when the data size is relatively small.

We did some additional tests to study the effect of the order in which the CD algorithm

cycles through pairs of nodes in steps 2 - 20 of Algorithm 1. Instead of a fixed order, we may use a random order each time we cycle through the edges. The results of these two implementations are largely comparable with no statistically significant differences. See supplemental material for the detailed comparisons.

5.3 Results for observational data

In this subsection, we apply our CD algorithm to observational data and compare its performance with the PC algorithm and the HC algorithm. For observational data, we have tested the algorithms on both low-dimensional data and high-dimensional data. Since we are more interested in the high-dimensional case, we will only report the detailed results for tests in this case. A brief summary of the test results on low-dimensional data will be provided, with the detailed data in the supplementary material. Since the direction of an edge may not be identifiable from observational data, we only compare estimated skeletons.

For observational data, we generated three graphs, with $p = 50, 100, 200$, for each type of the four networks. We chose $n \in \{p, 5p\}$ for low-dimensional cases and $(n, p) \in \{(50, 100), (50, 200)\}$ for high-dimensional cases. The number of true edges s_0 ranges from 49 to 400 for these graphs. Again, 20 data sets were generated for each graph.

For low-dimensional data, the HC algorithm predicted too many edges that the number of false positive edges was always large. As a result, this method always has the highest TPR, but also the highest FDR. Our CD algorithm predicted fewer edges, giving a lower TPR and a much lower FDR. Compared with the PC algorithm, like for interventional data, the CD algorithm has a much better performance for bipartite graphs and scale-free networks. For the other two networks, our method shows a comparable performance with the PC algorithm when $n = p$ and performs slightly worse when $n = 5p$. These test results on low-dimensional data are in good agreement with the conclusion drawn from interventional data that our CD algorithm has an advantage over the PC algorithm for a small sample size.

The problem of high false positive rates became even more severe when the HC algorithm was applied to the high-dimensional data sets. For example, for a scale-free network with $p = 200$, the HC algorithm predicted more than 2,000 edges. This makes its results not comparable

Table 4: Average results for skeleton estimation from high-dimensional observational data

Network	(n, p, s_0)	Method	P	E	FP	TPR	FDR
Bipartite	$(50, 100, 100)$	CD	73.2	57.8	15.4	0.578(0.049)	0.210(0.040)
		CD*	52.4	45.4	7.0	0.454(0.032)	0.134(0.044)
		PC	52.4	38.3	14.2	0.383(0.028)	0.270(0.045)
Polytree	$(50, 100, 99)$	CD	105.4	85.0	20.4	0.859(0.038)	0.193(0.041)
		CD*	68.6	63.2	5.4	0.639(0.039)	0.078(0.031)
		PC	68.8	63.6	5.2	0.642(0.025)	0.076(0.024)
Scale-free	$(50, 100, 99)$	CD	105.8	73.0	32.8	0.737(0.038)	0.310(0.033)
		CD*	62.4	53.5	8.9	0.540(0.029)	0.143(0.039)
		PC	62.2	43.0	19.1	0.435(0.029)	0.307(0.062)
Small-world	$(50, 100, 200)$	CD	95.8	87.9	7.9	0.440(0.018)	0.082(0.028)
		CD*	72.0	69.0	3.0	0.345(0.022)	0.041(0.019)
		PC	71.4	66.8	4.6	0.334(0.020)	0.064(0.025)
Bipartite	$(50, 200, 200)$	CD	124.1	96.4	27.8	0.482(0.034)	0.224(0.047)
		CD*	115.4	92.7	22.7	0.464(0.026)	0.197(0.031)
		PC	116.8	78.2	38.6	0.391(0.021)	0.331(0.027)
Polytree	$(50, 200, 199)$	CD	208.9	169.6	39.3	0.852(0.023)	0.188(0.016)
		CD*	136.6	126.0	10.6	0.633(0.023)	0.078(0.019)
		PC	137.9	124.4	13.5	0.625(0.020)	0.098(0.018)
Scale-free	$(50, 200, 199)$	CD	216.9	130.2	86.7	0.654(0.034)	0.400(0.037)
		CD*	131.0	103.0	28.0	0.517(0.029)	0.214(0.036)
		PC	131.0	80.4	50.6	0.404(0.021)	0.386(0.027)
Small-world	$(50, 200, 400)$	CD	226.0	156.0	70.0	0.390(0.020)	0.310(0.026)
		CD*	101.9	81.4	20.5	0.204(0.015)	0.201(0.039)
		PC	102.5	85.0	17.6	0.212(0.013)	0.171(0.032)

to those from the other methods. Therefore, we only compare our CD algorithm with the PC algorithm for high-dimensional data. Table 4 summarizes the results of these two methods. One sees that our model selection criterion works quite well: The number of predicted edges is close to or smaller than s_0 for most of the graphs, resulting in a good control of the FDR. In many cases, the DAG selected by our criterion has a higher TPR and a comparable or even lower FDR than the one estimated by the PC algorithm. The comparison becomes clearer after matching edges. For six out of the eight graphs, our CD algorithm shows simultaneously a higher TPR and a lower FDR. For the polytree with $p = 100$, the two methods have almost identical error rates. The only case in which the PC algorithm did a better job is the small-world network of

$p = 200$ nodes, but the differences are small.

The simulation results on both interventional and observational data demonstrate the advantages of our method when $n \leq p$ and when p is large. In these cases, our CD algorithm in general shows better (or comparable) performance than the PC algorithm. The EE-DAG sampler cannot be applied to graphs with $p \geq 100$ due to its demanding computation, while a DAG estimated by the HC algorithm is too dense with an overwhelming portion of false positive edges. The score function in the HC algorithm is a Bayesian information criterion under the product multinomial model (Section 2.1). Its inferior performance is likely due to overfitting to a model with too many free parameters, particularly in the high-dimensional scenario with limited amount of data. This comparison thus also highlights the advantage of a more parsimonious multi-logit model in our penalized likelihood approach.

5.4 Timing comparison

We comment briefly on the comparison of running time between our method and the PC algorithm. The running time for generating a whole solution path of 30 graphs for the interventional data (results in Tables 1 and 2) was between 0.8 and 6 minutes for most graphs, except the four with $n = 500$ and $p = 100$ for which the running time was around 30 minutes. The PC algorithm was faster, spending at most 2 minutes in all the cases. For high-dimensional data (Table 4), it took between 2 and 14 minutes for our method to compute the entire solution path, while the runtime for the PC algorithm was under one minute for all the data sets. Although the per-graph runtime of the CD algorithm is quite comparable, the PC algorithm has a clear advantage in terms of the overall speed. However, our method gives a principled way to incorporate interventional data and is often more accurate for data sets with small n and large p. These merits in performance justify its utility. In addition, there is room for a more efficient implementation of our algorithm which may improve its speed substantially, such as the use of stochastic gradient descent to save the computation of the gradient of the log-likelihood (13).

6 A Real Data Example

We consider in this section applying the CD algorithm to a real data set that has been extensively studied. The data set was generated from a flow cytometry experiment conducted by Sachs et al. (2005), who studied a well-known signaling network in human primary CD4+ T-cells of the immune system. This chosen network was perturbed by various stimulatory and inhibitory interventions. Each interventional condition was applied to an individual component of the network. Simultaneous measurements were taken on $p = 11$ proteins and phospholipids of this network from individual cells under each condition. Since three interventions were targeted at proteins that were not measured, samples collected under these conditions were observational. Among the 11 measured components, five proteins and phospholipids were perturbed. The data set contains measurements for $n = 5,400$ cells. Each variable has three levels (high, medium and low), and consequently, the size of a component group of $\boldsymbol{\beta}$ is 6 for this data set.

Figure 2A shows the known causal interactions among the 11 measured components of this signaling network. These causal relationships are well-established, and no consensus has been reached on interactions beyond those present in the network. This network structure is often used as the benchmark to assess the accuracy of an estimated network. Therefore, we call it the consensus model. Our estimated network by the CD algorithm with 15 edges is shown in Figure 2B. The DAG is qualitatively close to the consensus model. The detailed performance measures are reported in Table 5. As a comparison, we include the DAGs estimated by four competing methods, the order-graph sampler (Ellis and Wong 2008), the multi-domain sampler (Zhou 2011), the EE-DAG sampler, and the PC algorithm. Our CD algorithm shows a very competitive performance, predicting more expected edges and fewer reversed edges than the PC algorithm, the order-graph sampler, and the EE-DAG sampler did. Our method also has the least number of false positive edges among all the methods. The DAG estimated by the multi-domain sampler has a higher TPR but a higher FDR as well. The result of the PC algorithm on this data set is unsatisfactory. Its estimated DAG has too many missing and false positive edges, indicating that the PC algorithm did not accurately predict the skeleton.

Markov chain Monte Carlo (MCMC) methods for DAG estimation often have good perfor-

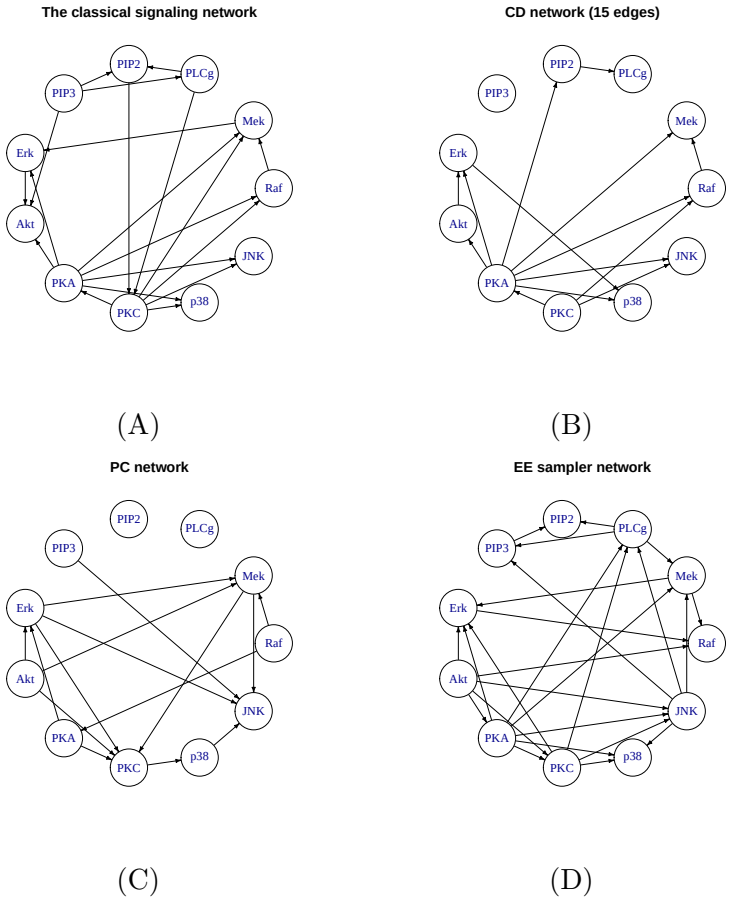

Figure 2: (A) The consensus signaling network in human immune system cells, (B) DAG estimated by the CD algorithm, (C) DAG estimated by the PC algorithm, (D) DAG estimated by the EE-DAG sampler.

mance when the number of nodes p is small, but they do not scale well. Thus, it is comforting to see that our method, which can handle larger networks, shows very competitive performance compared to MCMC methods on this relatively small network.

7 Discussion

We have developed a maximum penalized likelihood method for estimating sparse discrete Bayesian networks under a multi-logit model. In order to avoid penalizing separately individual dummy variables for a factor, the group Lasso penalty is utilized to encourage sparsity at the factor level. A blockwise coordinate descent algorithm is developed for parameter estimation

Table 5: Comparison on the flow cytometry data set

Method	P	E	R	M	FP
CD algorithm	15	10	2	8	3
PC algorithm	15	3	5	12	7
Order-graph sampler	20	8	4	8	8
EE-DAG sampler	26	9	6	5	11
Multi-domain sampler	25.9	15.55	2.05	2.4	8.3

Note: The order-graph sampler result comes from the mean graph (Figure 11 in Ellis and Wong 2008), while the multi-domain sampler result is the average over 20 independent runs (Table 3 in Zhou 2011).

where each coordinate descent step is solved by iteratively applying a quadratic approximation. The acyclicity constraint imposed on the structure of Bayesian networks can be enforced in a natural way by the coordinate descent algorithm. We have established finite-dimensional consistency for the proposed estimation procedure. Our method has been evaluated on four types of simulated graphs that resemble many real-world networks and has been tested on both interventional and observational data. We demonstrate that the performance of the CD algorithm is very competitive comparing to many existing methods, particularly when $n \leq p$. We have also performed an analysis of a flow cytometry data set generated from a signaling network in human immune system cells. The DAG estimated by the CD algorithm is close to the consensus model. Since the true network is not available, the estimated edges provide candidate causal interactions that could be tested in future experiments.

Computation for estimating discrete Bayesian networks is demanding due to the size of the parameter space and the nonlinear nature of the multi-logit model. There is room for improving the efficiency of the CD algorithm. For example, one may incorporate the idea of stochastic gradient in the quadratic approximation step, which will reduce significantly the computation. Moreover, since our search space is nonconvex, introducing such components of stochastic optimization may also increase the chance of finding the global minimizer of the penalized loss function. Other future directions include studying the consistency of our penalized estimator when the number of nodes $p = p_n$ grows with the sample size n and investigating the

use of group concave penalties.

Appendix: Proofs

We first review the definition of a topological sort of a DAG.

Definition 2 (Topological sorts)**.** A topological sort of a DAG $\mathcal{G}$ is a linear ordering $\sqsubset$ of its nodes such that $i \prec j$ in $\sqsubset$ if $i \in \Pi_j^{\mathcal{G}}$.

Due to its directed acyclic nature, every DAG has at least one topological sort. If an ordering $\sqsubset$ is a topological sort of a DAG $\mathcal{G}$, we say that $\mathcal{G}$ and $\sqsubset$ are compatible. Let $\mathcal{S}(\mathcal{G})$ denote the set of all topological sorts of a DAG $\mathcal{G}$.

Lemma A.1. *For any $\boldsymbol{\phi}^0 \in \Omega$, there is a $\delta(\boldsymbol{\phi}^0) > 0$ such that if $\boldsymbol{\phi} \in \Omega$ and $\|\boldsymbol{\phi} - \boldsymbol{\phi}^0\|_2 < \delta(\boldsymbol{\phi}^0)$ then $\mathcal{S}(\mathcal{G}_{\boldsymbol{\phi}}) \cap \mathcal{S}(\mathcal{G}_{\boldsymbol{\phi}^0}) \neq \varnothing$.*

Proof. If $\mathcal{G}_{\boldsymbol{\phi}^0}$ is an empty graph which is compatible with any ordering, the statement holds trivially. Otherwise, let $\delta(\boldsymbol{\phi}^0) = \frac{1}{2}\min_{j:\phi_j^0 \neq \mathbf{0}} \|\phi_j^0\|_2 > 0$, where ϕ_j^0 is the j^{th} component group of $\boldsymbol{\phi}^0$ (see Section 4). If $\|\boldsymbol{\phi} - \boldsymbol{\phi}^0\|_2 < \delta(\boldsymbol{\phi}^0)$ and $\phi_i^0 \neq \mathbf{0}$ for some $i \in \{1, \ldots, p^2\}$, then $\phi_i \neq \mathbf{0}$ as well, since otherwise $\|\boldsymbol{\phi} - \boldsymbol{\phi}^0\|_2 \geq \|\phi_i^0\|_2 > \delta(\boldsymbol{\phi}^0)$. This implies that every edge in $\mathcal{G}_{\boldsymbol{\phi}^0}$ is also an edge in $\mathcal{G}_{\boldsymbol{\phi}}$ and thus $\mathcal{G}_{\boldsymbol{\phi}}$ and $\mathcal{G}_{\boldsymbol{\phi}^0}$ have at least one common topological sort. □

This lemma shows that every DAG in a sufficiently small neighborhood of $\boldsymbol{\phi}^0$, denoted by $\text{nb}(\boldsymbol{\phi}^0) \subset \Omega$, has a topological sort that is also compatible with $\mathcal{G}_{\boldsymbol{\phi}^0}$.

Proof of Theorem 1. We prove the first claim by contradiction. Suppose $\boldsymbol{\phi} \neq \boldsymbol{\phi}^*$ and $p(\mathbf{x}|\boldsymbol{\phi}_{[k]}) = p(\mathbf{x}|\boldsymbol{\phi}^*_{[k]})$ for $k = 1, \ldots, p$. There are two cases to consider depending on the topological sorts of $\mathcal{G}_{\boldsymbol{\phi}}$ and those of $\mathcal{G}_{\boldsymbol{\phi}^*}$.

Case 1: $\mathcal{S}(\mathcal{G}_{\boldsymbol{\phi}}) \cap \mathcal{S}(\mathcal{G}_{\boldsymbol{\phi}^*}) \neq \varnothing$. Let $\sqsubset \in \mathcal{S}(\mathcal{G}_{\boldsymbol{\phi}}) \cap \mathcal{S}(\mathcal{G}_{\boldsymbol{\phi}^*})$, i.e., an ordering compatible with both $\mathcal{G}_{\boldsymbol{\phi}}$ and $\mathcal{G}_{\boldsymbol{\phi}^*}$. Assume without loss of generality that in this ordering $i \prec j$ if $i < j$. Apparently, $\sqsubset$ is also compatible with $\mathcal{G}_{\boldsymbol{\phi}_{[k]}}$ and $\mathcal{G}_{\boldsymbol{\phi}^*_{[k]}}$ for $k = 1, \ldots, p$. Then we can write $p(\mathbf{x}|\boldsymbol{\phi}_{[k]}) = \prod_{i=1}^p p(x_i|x_1, \ldots, x_{i-1}, \boldsymbol{\phi}_{[k]}) = \prod_{i=1}^p p(x_i|\Pi_i^{\mathcal{G}_{\boldsymbol{\phi}_{[k]}}}, \boldsymbol{\phi}_{[k]})$ and $p(\mathbf{x}|\boldsymbol{\phi}^*_{[k]}) =$

$\prod_{i=1}^{p} p(x_i|x_1,\ldots,x_{i-1},\boldsymbol{\phi}^*_{[k]}) = \prod_{i=1}^{p} p(x_i|\Pi_i^{\mathcal{G}_{\boldsymbol{\phi}^*_{[k]}}},\boldsymbol{\phi}^*_{[k]})$. Since $p(\mathbf{x}|\boldsymbol{\phi}_{[k]}) = p(\mathbf{x}|\boldsymbol{\phi}^*_{[k]})$ for all possible $\mathbf{x}$, it follows that $\Pi_i^{\mathcal{G}_{\boldsymbol{\phi}_{[k]}}} = \Pi_i^{\mathcal{G}_{\boldsymbol{\phi}^*_{[k]}}}$ for all i and thus $\mathcal{G}_{\boldsymbol{\phi}_{[k]}} = \mathcal{G}_{\boldsymbol{\phi}^*_{[k]}}$ for all k. However, since $\boldsymbol{\phi} \neq \boldsymbol{\phi}^*$, there exists some k such that $\boldsymbol{\phi}_{[k]} \neq \boldsymbol{\phi}^*_{[k]}$. Therefore, there exists a k such that the common probability mass function $p(\mathbf{x}|\boldsymbol{\phi}_{[k]}) = p(\mathbf{x}|\boldsymbol{\phi}^*_{[k]})$, factorized according to a common structure $\mathcal{G}_{\boldsymbol{\phi}_{[k]}} = \mathcal{G}_{\boldsymbol{\phi}^*_{[k]}}$, can be parameterized by two different parameters $\boldsymbol{\phi}_{[k]}$ and $\boldsymbol{\phi}^*_{[k]}$. This is impossible since the multi-logit model (4), subject to constraints (5) and (6), for $[X_j|\Pi_j^{\mathcal{G}_{\boldsymbol{\phi}}}]$ is identifiable.

Case 2: $\mathcal{S}(\mathcal{G}_{\boldsymbol{\phi}}) \cap \mathcal{S}(\mathcal{G}_{\boldsymbol{\phi}^*}) = \varnothing$, that is, none of the orderings of $\mathcal{G}_{\boldsymbol{\phi}^*}$ is compatible with $\mathcal{G}_{\boldsymbol{\phi}}$. In this case, there must exist a pair of indices (i,j) such that in $\mathcal{G}_{\boldsymbol{\phi}^*}$ $X_i \in \text{an}(X_j)$, but in $\mathcal{G}_{\boldsymbol{\phi}}$ X_j is a non-descendant of X_i. Then X_j is independent of X_i in the DAG of $\boldsymbol{\phi}_{[i]}$, since in $\mathcal{G}_{\boldsymbol{\phi}_{[i]}}$ X_i has no parents and X_j is a non-descendant of X_i. However, in $\mathcal{G}_{\boldsymbol{\phi}^*_{[i]}}$ we still have $X_i \in \text{an}(X_j)$. Since $\boldsymbol{\phi}^*$ is natural, X_i and X_j are not independent in $\boldsymbol{\phi}^*_{[i]}$. Therefore, there exists $1 \leq i \leq p$ such that $p(\mathbf{x}|\boldsymbol{\phi}_{[i]}) \neq p(\mathbf{x}|\boldsymbol{\phi}^*_{[i]})$, which contradicts our assumption.

So in both *case 1* and *case 2* we have a contradiction. Thus, the first claim holds.

To prove the second claim (24), note that by the law of large numbers,

$$\begin{aligned}\frac{1}{n}(L(\boldsymbol{\phi}) - L(\boldsymbol{\phi}^*)) &= \sum_{k=1}^{p} \frac{n_k}{n} \frac{1}{n_k} \sum_{h \in \mathcal{I}_k} \log \frac{p(\mathcal{X}_{h\cdot}|\boldsymbol{\phi}_{[k]})}{p(\mathcal{X}_{h\cdot}|\boldsymbol{\phi}^*_{[k]})} \\ &\to_p \sum_{k=1}^{p} \alpha_k \mathbf{E}_{\boldsymbol{\phi}^*_{[k]}} \left[\log \frac{p(\mathbf{Y}|\boldsymbol{\phi}_{[k]})}{p(\mathbf{Y}|\boldsymbol{\phi}^*_{[k]})} \right], \end{aligned} \tag{A.1}$$

where $\mathbf{Y}$ is a random vector with probability mass function $p(\mathbf{x}|\boldsymbol{\phi}^*_{[k]})$. Then the desired result follows immediately using Jensen's inequality and the first claim proved above. □

Proof of Theorem 2. Let $\mathbf{Y}$ be a random vector with probability mass function $p(\mathbf{x}|\boldsymbol{\phi})$ and

$$\mathbf{I}(\boldsymbol{\phi}) = \mathbf{E}_{\boldsymbol{\phi}} \left\{ \left[\frac{\partial}{\partial \boldsymbol{\phi}} \log p(\mathbf{Y}|\boldsymbol{\phi}) \right] \left[\frac{\partial}{\partial \boldsymbol{\phi}} \log p(\mathbf{Y}|\boldsymbol{\phi}) \right]^T \right\}$$

be the Fisher information matrix. We consider two cases.

Case 1: Consider $\boldsymbol{\phi} \in \text{nb}(\boldsymbol{\phi}^*)$. By Lemma A.1, $\mathcal{G}_{\boldsymbol{\phi}}$ and $\mathcal{G}_{\boldsymbol{\phi}^*}$ have a common compatible ordering. If we restrict to the lower dimensional space $\Omega_{[k]} = \{\boldsymbol{\phi}_{[k]} : \boldsymbol{\phi} \in \Omega\}$, the same argument

applies to an arbitrarily small neighborhood of $\boldsymbol{\phi}^*_{[k]}$ in this space, that is, $\mathcal{G}_{\boldsymbol{\phi}_{[k]}}$ and $\mathcal{G}_{\boldsymbol{\phi}^*_{[k]}}$ share a compatible ordering. Then it follows from the arguments used in *Case 1* in the proof of Theorem 1 that, for $\boldsymbol{\phi}_{[k]} \in \text{nb}(\boldsymbol{\phi}^*_{[k]}) \backslash \{\boldsymbol{\phi}^*_{[k]}\}$, $p(\mathbf{x}|\boldsymbol{\phi}_{[k]}) \neq p(\mathbf{x}|\boldsymbol{\phi}^*_{[k]})$ for some non-negligible set of $\mathbf{x}$ and

$$\mathbf{E}_{\boldsymbol{\phi}^*_{[k]}}\left[\log p(\mathbf{Y}|\boldsymbol{\phi}^*_{[k]})\right] > \mathbf{E}_{\boldsymbol{\phi}^*_{[k]}}\left[\log p(\mathbf{Y}|\boldsymbol{\phi}_{[k]})\right].$$

This implies that $\mathbf{I}(\boldsymbol{\phi}^*_{[k]})$ is positive definite for all k.

Let $\mathbf{u} \in \{\mathbf{u} : \boldsymbol{\phi}^* + a_n\mathbf{u} \in \Omega\}$ and u_j be its j^{th} component group defined in the same way as ϕ_j. Without loss of generality, we can assume that $\|\mathbf{u}\|_2 = M$ where M is a fixed positive constant. Further, let $\mathbf{u}_{[k]}$ be the vector defined similarly as $\boldsymbol{\phi}_{[k]}$ by setting $u_j = \mathbf{0}$ if the j^{th} group corresponds to an edge pointing to X_k. Note that $\sum_{k=1}^p \|\mathbf{u}_{[k]}\|_2^2 \geq \|\mathbf{u}\|_2^2$. Let $\delta^k_{\min} > 0$ be the minimal eigenvalue of $\mathbf{I}(\boldsymbol{\phi}^*_{[k]})$ and $\rho = \min_k(\alpha_k\delta^k_{\min}/2)$. Then

$$\sum_{k=1}^p \frac{\alpha_k}{2}\mathbf{u}^T_{[k]}\mathbf{I}(\boldsymbol{\phi}^*_{[k]})\mathbf{u}_{[k]} \geq \sum_{k=1}^p \frac{\alpha_k}{2}\delta^k_{\min}\|\mathbf{u}_{[k]}\|_2^2 \geq \rho\sum_{k=1}^p\|\mathbf{u}_{[k]}\|_2^2 \geq \rho\|\mathbf{u}\|_2^2. \tag{A.2}$$

Write $a_n = \omega(b_n)$ if $b_n^{-1} = O(a_n^{-1})$. Let $\omega(1/\sqrt{n}) = a_n = o(1)$. Recall that $\mathcal{B} = \{j : \phi^*_j \neq \mathbf{0}, 1 \leq j \leq p^2\}$. We have

$$\begin{aligned}
&R(\boldsymbol{\phi}^* + a_n\mathbf{u}) - R(\boldsymbol{\phi}^*) \\
\leq\ &L(\boldsymbol{\phi}^* + a_n\mathbf{u}) - L(\boldsymbol{\phi}^*) - \lambda_n\sum_{j\in\mathcal{B}}\tau_j(\|\phi^*_j + a_nu_j\|_2 - \|\phi^*_j\|_2) \\
\leq\ &\sum_{k=1}^p\left[L_k(\boldsymbol{\phi}^*_{[k]} + a_n\mathbf{u}_{[k]}) - L_k(\boldsymbol{\phi}^*_{[k]})\right] + \lambda_n a_n\sum_{j\in\mathcal{B}}\tau_j\|u_j\|_2,
\end{aligned} \tag{A.3}$$

where the last line follows from the triangle inequality. Taylor expansion of L_k around $\boldsymbol{\phi}^*_{[k]}$ in

(A.3) leads to

$$\begin{aligned}
&\sum_{k=1}^{p}\left[a_n\{\nabla L_k(\boldsymbol{\phi}^*_{[k]})\}^T\mathbf{u}_{[k]}-\frac{1}{2}n_k a_n^2\mathbf{u}_{[k]}^T\mathbf{I}(\boldsymbol{\phi}^*_{[k]})\mathbf{u}_{[k]}\left\{1+o_p(1)\right\}\right]+\lambda_n a_n\sum_{j\in\mathcal{B}}\tau_j\|u_j\|_2\\
\leq&\sum_{k=1}^{p}\left[\sqrt{\alpha_k}\sqrt{n}a_n\frac{\{\nabla L_k(\boldsymbol{\phi}^*_{[k]})\}^T}{\sqrt{n_k}}\mathbf{u}_{[k]}\left\{1+o_p(1)\right\}\right]-\rho n a_n^2\|\mathbf{u}\|_2^2\left\{1+o_p(1)\right\}\\
&+\lambda_n n^{-1/2}\sqrt{n}a_n\sum_{j\in\mathcal{B}}\tau_j\|u_j\|_2,
\end{aligned}$$

where we have used (A.2) and that $n_k/n\to\alpha_k$. Now dividing both sides by $\sqrt{n}a_n$, we arrive at

$$\begin{aligned}
&(\sqrt{n}a_n)^{-1}\left[R(\boldsymbol{\phi}^*+a_n\mathbf{u})-R(\boldsymbol{\phi}^*)\right]\\
\leq&\sum_{k=1}^{p}\left[\sqrt{\alpha_k}\frac{\{\nabla L_k(\boldsymbol{\phi}^*_{[k]})\}^T}{\sqrt{n_k}}\mathbf{u}_{[k]}\left\{1+o_p(1)\right\}\right]-\rho\sqrt{n}a_n\|\mathbf{u}\|_2^2\left\{1+o_p(1)\right\}\\
&+\lambda_n n^{-1/2}\sum_{j\in\mathcal{B}}\tau_j\|u_j\|_2. \qquad\text{(A.4)}
\end{aligned}$$

Note that $n_k^{-1/2}\|\nabla L_k(\boldsymbol{\phi}^*_{[k]})\|_2=O_p(1)$ for all k by the central limit theorem, $\tau_j=O_p(1)$ for $j=1,\ldots,p^2$, and $\lambda_n/\sqrt{n}=o(1)$. Therefore, the second order term in (A.4) dominates the first and the third terms uniformly if $\sqrt{n}a_n$ is sufficiently large. Hence, for any $\varepsilon>0$, there exists a constant $C<\infty$ such that

$$P\left[\sup_{a_n\|\mathbf{u}\|_2\geq C/\sqrt{n}}R(\boldsymbol{\phi}^*+a_n\mathbf{u})<R(\boldsymbol{\phi}^*)\right]\geq 1-\varepsilon, \qquad\text{(A.5)}$$

when n is sufficiently large.

Case 2: Consider $\boldsymbol{\phi}\in\Omega\backslash\mathrm{nb}(\boldsymbol{\phi}^*)$. Let $\sqsubset_i$ $(1\leq i\leq p!)$ be an ordering of the p variables. Since $R(\boldsymbol{\phi})$ is a strictly concave function in the subspace $\Omega_i=\{\boldsymbol{\phi}\in\Omega:\sqsubset_i \text{ is compatible with } \mathcal{G}_{\boldsymbol{\phi}}\}$, there exists a unique local maximizer $\hat{\boldsymbol{\phi}}^i$ of $R(\boldsymbol{\phi})$ in Ω_i. The pointwise convergence in probability of the random concave function $n^{-1}R(\boldsymbol{\phi})$ implies that $\hat{\boldsymbol{\phi}}^i\to_p\boldsymbol{\phi}^i$, where $\boldsymbol{\phi}^i\in\Omega_i$ is the unique maximizer of its limiting function; see Andersen and Gill (1982) and Pollard (1991). Let

$\mathcal{M} \triangleq \{\boldsymbol{\phi}^i : \boldsymbol{\phi}^i \neq \boldsymbol{\phi}^*\}$ which is a finite set. For any $\boldsymbol{\phi}^i \in \mathcal{M}$ we have

$$\begin{aligned}\frac{1}{n}R(\boldsymbol{\phi}^i) - \frac{1}{n}R(\boldsymbol{\phi}^*) &\leq \frac{1}{n}L(\boldsymbol{\phi}^i) - \frac{1}{n}L(\boldsymbol{\phi}^*) - \frac{1}{n}\lambda_n \sum_{j\in\mathcal{B}} \tau_j(\|\phi_j^i\|_2 - \|\phi_j^*\|_2) \\ &\to_p \sum_{k=1}^{p} \alpha_k \mathbf{E}_{\boldsymbol{\phi}^*_{[k]}}\left[\log \frac{p(\mathbf{Y}|\boldsymbol{\phi}^i_{[k]})}{p(\mathbf{Y}|\boldsymbol{\phi}^*_{[k]})}\right] - \frac{1}{n}\lambda_n O(1), \end{aligned} \tag{A.6}$$

by (A.1). The second term is $o(\frac{1}{\sqrt{n}})$ since $\lambda_n = o(\sqrt{n})$, while the first term is a negative constant by Theorem 1. Thus, the last line converges to a negative constant. Since $|\mathcal{M}|$ is finite, the convergence is uniform over all $\boldsymbol{\phi}^i \in \mathcal{M}$. Therefore, for any $\varepsilon > 0$,

$$P\left[\sup_{\mathcal{M}} R(\boldsymbol{\phi}^i) < R(\boldsymbol{\phi}^*)\right] \geq 1 - \varepsilon, \tag{A.7}$$

when n is sufficiently large.

Combining (A.5) and (A.7), we have shown that there is a global maximizer $\hat{\boldsymbol{\phi}}$ of $R(\boldsymbol{\phi})$ such that $\|\hat{\boldsymbol{\phi}} - \boldsymbol{\phi}^*\|_2 = O_p(n^{-1/2})$. □

Proof of Theorem 3. Under our choice of the weights, $\tau_j = O_p(1)$ for all $j \in \mathcal{B}$, since $\tilde{\phi}_j$ is consistent for ϕ_j^*. By a similar proof to that of Theorem 2, one can show that $\|\hat{\boldsymbol{\phi}} - \boldsymbol{\phi}^*\|_2 = O_p(n^{-1/2})$, where $\hat{\boldsymbol{\phi}}$ is a global maximizer.

Recall that $\mathcal{A} = \{j : \phi_j^* = \mathbf{0}, 1 \leq j \leq p^2\}$. To prove the second claim, let us first, by permuting the indices, rewrite the parameter $\boldsymbol{\phi}$ as $\boldsymbol{\phi} = (\boldsymbol{\phi}_a, \boldsymbol{\phi}_b) = (\boldsymbol{\phi}_{\mathcal{A}}, \boldsymbol{\phi}_{\mathcal{B}}, \boldsymbol{\phi}_{(2)})$, where $\boldsymbol{\phi}_a = \boldsymbol{\phi}_{\mathcal{A}}$ and $\boldsymbol{\phi}_b = (\boldsymbol{\phi}_{\mathcal{B}}, \boldsymbol{\phi}_{(2)})$. Now we only need to show that with probability tending to 1, for any $\boldsymbol{\phi}_b$ satisfying $\|\boldsymbol{\phi}_b - \boldsymbol{\phi}_b^*\|_2 = O(n^{-1/2})$ and any constant $C > 0$,

$$\operatorname*{arg\,max}_{\boldsymbol{\phi}_a : \|\boldsymbol{\phi}_a\|_2 \leq C/\sqrt{n}} R\big((\boldsymbol{\phi}_a, \boldsymbol{\phi}_b)\big) = \mathbf{0}. \tag{A.8}$$

To establish (A.8), we again study the behavior of $R(\boldsymbol{\phi})$ around the point $\boldsymbol{\phi}^o = (\mathbf{0}, \boldsymbol{\phi}_b)$ by expanding $L(\boldsymbol{\phi})$. Let $\mathbf{u} = (\mathbf{u}_a, \mathbf{u}_b)$ such that $\|\mathbf{u}_a\|_2 \leq C$, $\mathbf{u}_b = \mathbf{0}$, and $\boldsymbol{\phi}^o + n^{-1/2}\mathbf{u} \in \Omega$. Let $\mathcal{A}_0 = \{j \in \mathcal{A} : \phi^*_{r(j)} = \mathbf{0}\}$, where $r(j)$ is defined in (25). Since any DAG cannot simultaneously have a pair of reserved edges ($i \to j$ and $j \to i$), we must have $u_j = \mathbf{0}$ for any $j \in \mathcal{A} \setminus \mathcal{A}_0$. Then

we have the following result similar to that in the proof of Theorem 2:

$$
\begin{aligned}
& R(\boldsymbol{\phi}^o + n^{-1/2}\mathbf{u}) - R(\boldsymbol{\phi}^o) \\
= & \sum_{k=1}^{p} \left[\sqrt{\alpha_k} \frac{\{\nabla L_k(\boldsymbol{\phi}^o_{[k]})\}^T}{\sqrt{n_k}} \mathbf{u}_{[k]} \{1 + o_p(1)\} - \frac{\alpha_k}{2} \mathbf{u}^T_{[k]} \mathbf{I}(\boldsymbol{\phi}^o_{[k]}) \mathbf{u}_{[k]} \{1 + o_p(1)\} \right] \\
& - \frac{\lambda_n}{\sqrt{n}} n^{\gamma/2} \sum_{j \in \mathcal{A}_0} \|\sqrt{n}\tilde{\phi}_j\|_2^{-\gamma} \|u_j\|_2. \qquad \text{(A.9)}
\end{aligned}
$$

Note that both the first and the second terms on the right-hand side of (A.9) are on the order of $O_p(1)$ for any fixed constant C. Since $\tilde{\phi}_j$ is $\sqrt{n}$-consistent, $\|\sqrt{n}\tilde{\phi}_j\|_2 = O_p(1)$ for $j \in \mathcal{A}_0$ and the third term is on the order of $\lambda_n n^{(\gamma-1)/2} \to \infty$. Therefore, (A.8) holds, and the proof is complete. □

References

Andersen, P. K. and Gill, R. D. (1982), "Cox's regression model for counting processes: a large sample study," *The Annals of Statistics*, 10, 1100–1120.

Aragam, B. and Zhou, Q. (2015), "Concave penalized estimation of sparse Bayesian networks," *Journal of Machine Learning Research*, in press.

Bach, F. R. (2008), "Consistency of the Group Lasso and Multiple Kernel Learning," *Journal of Machine Learning Research*, 9, 1179–1225.

Barabási, A.-L. and Albert, R. (1999), "Emergence of Scaling in Random Networks," *Science*, 286, 509–512.

Bouckaert, R. R. (1993), "Probabilistic Network Construction Using the Minimum Description Length Principle," in *Symbolic and Quantitative Approaches to Reasoning and Uncertainty: European Conference ECSQARU '93*, Springer, vol. 747 of *Lecture Notes in Computer Science*, pp. 41–48.

— (1994), "Probabilistic Network Construction Using the Minimum Description Length Principle," Tech. Rep. RUU-CS-94-27, Department of Computer Science, Utrecht University.

Buntine, W. (1991), "Theory Refinement on Bayesian Networks," in *Proceedings of the Seventh Annual Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann, pp. 52–60.

Chickering, D. M. and Heckerman, D. (1997), "Efficient Approximations for the Marginal Likelihood of Bayesian Networks with Hidden Variables," *Machine Learning*, 29, 181–212.

Cooper, G. F. and Herskovits, E. (1992), "A Bayesian Method for the Induction of Probabilistic Networks from Data," *Machine Learning*, 9, 309–347.

Cooper, G. F. and Yoo, C. (1999), "Causal discovery from a mixture of experimental and observational data," in *Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence*, Morgan Kaufmann Publishers Inc., pp. 116–125.

Croissant, Y. (2011), "Estimation of Multinomial Logit Models in R: The mlogit Packages," Available at: http://cran.r-project.org/web/packages/mlogit/vignettes/mlogit.pdf.

Csárdi, G. and Nepusz, T. (2006), "The igraph Software Package for Complex Network Research," *Proceedings of the International Conference on Complex Systems*.

Ellis, B. and Wong, W. H. (2008), "Learning Causal Bayesian Network Structures from Experimental Data," *Journal of the American Statistical Association*, 103, 778–789.

Friedman, J., Hastie, T., Höfling, H., and Tibshirani, R. (2007), "Pathwise Coordinate Optimization," *The Annals of Applied Statistics*, 1, 302–332.

Friedman, J., Hastie, T., and Tibshirani, R. (2010), "Regularization Paths for Generalized Linear Models via Coordinate Descent," *Journal of Statistical Software*, 33, 1–22.

Fu, F. and Zhou, Q. (2013), "Learning Sparse Causal Gaussian Networks With Experimental Intervention: Regularization and Coordinate Descent," *Journal of the American Statistical Association*, 108, 288–300.

Fu, W. (1998), "Penalized Regressions: The Bridge versus the Lasso," *Journal of Computational and Graphical Statistics*, 7, 397–416.

Gámez, J. A., Mateo, J. L., and Puerta, J. M. (2011), "Learning Bayesian networks by hill climbing: efficient methods based on progressive restriction of the neighborhood," *Data Mining and Knowledge Discovery*, 22, 106–148.

Heckerman, D., Geiger, D., and Chickering, D. M. (1995), "Learning Bayesian Networks: The Combination of Knowledge and Statistical Data," *Machine Learning*, 20, 197–243.

Herskovits, E. and Cooper, G. (1990), "Kutató: An Entropy-Driven System for Construction of Probabilistic Expert Systems from Databases," in *Proceedings of the Sixth Annual Conference on Uncertainty in Artificial Intelligence*, pp. 54–62.

Kalisch, M. and Bühlmann, P. (2007), "Estimating high-dimensional directed acyclic graphs with the PC-algorithm," *The Journal of Machine Learning Research*, 8, 613–636.

Kalisch, M., Mächler, M., Colombo, D., Maathuis, M. H., and Bühlmann, P. (2012), "Causal inference using graphical models with the R package pcalg," *Journal of Statistical Software*, 47, 1–26.

Kim, Y., Kim, J., and Kim, Y. (2006), "Blockwise Sparse Regression," *Statistica Sinica*, 16, 375–390.

Kou, S., Zhou, Q., and Wong, W. H. (2006), "Equi-energy sampler with applications in statistical inference and statistical mechanics (with discussion)," *The Annals of Statistics*, 34, 1581–1652.

Lam, W. and Bacchus, F. (1994), "Learning Bayesian Belief Networks: An Approach Based on the MDL Principle," *Computational Intelligence*, 10, 269–293.

Meganck, S., Leray, P., and Manderick, B. (2006), "Learning causal bayesian networks from observations and experiments: A decision theoretic approach," in *Modeling Decisions for Artificial Intelligence*, Springer, pp. 58–69.

Meier, L., van de Geer, S., and Bühlmann, P. (2008), "The Group Lasso for Logistic Regression," *Journal of the Royal Statistical Society. Series B.*, 70, 53–71.

Pearl, J. (2000), *Causality: Models, Reasoning, and Inference*, Cambridge University Press.

Peér, D., Regev, A., Elidan, G., and Friedman, N. (2001), "Inferring subnetworks from perturbed expression profiles," *Bioinformatics*, 17, S215–S224.

Pollard, D. (1991), "Asymptotics for least absolute deviation regression estimators," *Econometric Theory*, 7, 186–199.

Pournara, I. and Wernisch, L. (2004), "Reconstruction of gene networks using Bayesian learning and manipulation experiments," *Bioinformatics*, 20, 2934–2942.

Renyi, A. and Erdos, P. (1959), "On random graphs," *Publicationes Mathematicae*, 6, 5.

Russell, S. and Norvig, P. (1995), "Artificial intelligence: a modern approach," .

Sachs, K., Perez, O., Pe'er, D., Lauffenburger, D. A., and Nolan, G. P. (2005), "Causal Protein-Signaling Networks Derived from Multiparameter Single-Cell Data," *Science*, 308, 523–529.

Schmidt, M., Niculescu-Mizil, A., Murphy, K., et al. (2007), "Learning graphical model structure using L1-regularization paths," in *AAAI*, vol. 7, pp. 1278–1283.

Scutari, M. (2010), "Learning Bayesian Networks with the bnlearn R Package," *Journal of Statistical Software*, 35(i03).

Shojaie, A. and Michailidis, G. (2010), "Penalized Likelihood Methods for Estimation of Sparse High-Dimensional Directed Acyclic Graphs," *Biometrika*, 97, 519–538.

Spirtes, P., Glymour, C., and Scheines, R. (1993), *Causation, Prediction, and Search*, Springer-Verlag.

Suzuki, J. (1993), "A Construction of Bayesian Networks from Databases Based on an MDL Principle," in *Proceedings of the Ninth Annual Conference on Uncertainty in Artificial Intelligence*, pp. 266–273.

Tsamardinos, I., Brown, L. E., and Aliferis, C. F. (2006), "The max-min hill-climbing Bayesian network structure learning algorithm," *Machine learning*, 65, 31–78.

Tseng, P. and Yun, S. (2009), "A Coordinate Gradient Descent Method for Nonsmooth Separable Minimization," *Mathematical Programming B*, 117, 387–423.

van de Geer, S. and Bühlmann, P. (2013), "ℓ_0-penalized maximum likelihood for sparse directed acyclic graphs," *The Annals of Statistics*, 41, 536–567.

Wang, H. and Leng, C. (2008), "A Note on Adaptive Group Lasso," *Computational Statistics and Data Analysis*, 52, 5277–5286.

Watts, D. J. and Strogatz, S. H. (1998), "Collective Dynamics of 'Small-World' Networks," *Nature*, 393, 440–442.

Wei, F. and Huang, J. (2010), "Consistent Group Selection in High-Dimensional Linear Regression," *Bernoulli*, 16, 1369–1384.

Wu, T. and Lange, K. (2008), "Coordinate Descent Procedures for Lasso Penalized Regression," *The Annals of Applied Statistics*, 2, 224–244.

Xiang, J. and Kim, S. (2013), "A* Lasso for learning a sparse Bayesian network structure for continuous variables," in *Advances in Neural Information Processing Systems*, pp. 2418–2426.

Yuan, M. and Lin, Y. (2006), "Model Selection and Estimation in Regression with Grouped Variables," *Journal of the Royal Statistical Society. Series B*, 68, 49–67.

Zhou, Q. (2011), "Multi-Domain Sampling with Applications to Structural Inference of Bayesian Networks," *Journal of the American Statistical Association*, 106, 1317–1330.

Zhu, J. and Hastie, T. (2004), "Classification of Gene Microarrays by Penalized Logistic Regression," *Biostatistics*, 5, 427–443.